

A Low Power, Low Noise, 1.8GHz Voltage-Controlled Oscillator

by

Donald A. Hitko

Bachelor of Science in Electrical Engineering
GMI Engineering & Management Institute, 1994

Submitted to the Department of Electrical Engineering
and Computer Science in partial fulfillment of the
requirements for the degree of

Master of Science in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 1997

© 1997 Massachusetts Institute of Technology
All rights reserved.

Author
Department of Electrical Engineering and Computer Science
September 27, 1996

Certified by
Charles G. Sodini
Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by
F. R. Morgenthaler
Chairman, Department Committee on Graduate Students

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

MAR 06 1997

LIBRARIES

ENG

A Low Power, Low Noise, 1.8GHz Voltage-Controlled Oscillator

by

Donald A. Hitko

Submitted to the Department of Electrical Engineering and
Computer Science on September 27, 1996, in partial fulfillment of
the requirements for the degree of Master of Science in
Electrical Engineering.

Abstract

Transceivers which form the core of many wireless communications products often require a low noise voltage-controlled oscillator (VCO) for frequency shifting or synthesis in RF front-end circuitry. Since many wireless applications are focused on portability, low power operation is a necessity. Techniques for implementing oscillators are explored and then evaluated for the purposes of simultaneously realizing low noise and low power operation. Models for the design of oscillators and the analysis of oscillation stability are covered, and methods of calculating phase noise are discussed.

These models and theories all point to the need for a high quality, passive, integrated inductor to meet the system goals. Results from an experimental study of spiral and bond wire inductors built onto a silicon substrate are presented. The information gleaned from this study was used in selecting an inductor built from bond wires for use in a VCO for 1.8GHz applications. Using characterization data, a model for the inductor was constructed for circuit simulation purposes.

The design of a differential VCO circuit in a silicon bipolar process is detailed, including transistor considerations, development of a low voltage topology, and noise matching to the oscillator gain stage. On-chip varactors are integrated as a vital component of a tunable resonator, and a number of interface issues which impact the VCO design are introduced. Circuit simulation results which demonstrate the robustness of the oscillator are provided. Preliminary measurements on the fabricated VCO circuits gauge the oscillator noise spectral density to be -118dBc/Hz at a 1 MHz offset, drawing just 5mA from a supply of only 1.8V. Oscillation is supported with power consumption levels as small as 3.2mW from a 0.91 V supply, achieving -96dBc/Hz at the same carrier offset.

Thesis Supervisor: Charles G. Sodini
Title: Professor of Electrical Engineering

Acknowledgments

For my initial foray into RF/microwave circuit design, I've found the outcome of this effort to be most satisfying. It has not been a trail on which I embarked alone, and support has been plentiful along the way. Guidance has been supplied by Charlie Sodini who, in supervising this project, has pushed me to find answers, and not just solutions. He is deserving of a badge of valor for reading and re-reading this thesis, and for continually challenging my grasp of both device physics and the English language.

Our industrial partner in this journey has proven to be invaluable, and numerous people in the fledgling Communications Division of Analog Devices, Inc. (ADI) in Wilmington, Massachusetts are owed a debt of gratitude. ADI performed the wafer fabrication and made available its RF laboratory and design expertise. A special note of thanks goes to Ted Tewksbury, who provided his time, technical assistance, and above all, made things happen.

Additional gratitude is extended to the staff of the Microsystems Technology Laboratories for the fabrication of test wafers, and particularly to Matt Lukaszewski, for wire bonding many of my crazy concoctions.

Appreciation is expressed to Delco Electronics Corporation, first in Flint, Michigan, and more recently the IC Design Center in Kokomo, Indiana. Their patience and willingness to work with me in new avenues (and venues!) of exploration are greatly valued. Being able to spend time with the people of the Microwave Circuits and Devices group—who seem imbued with a nearly limitless patience for me—within Hughes Research Laboratories in Malibu, California has been a major asset in my learning of this field.

Certainly deserving of a mention are my colleagues in Sodini Group, who have helped me get things done, and perhaps most importantly, have shown me how the academic world works.

Finally, I should end where it all began. To my parents, whose unwavering support throughout the years has been appreciated to a level that a mere acknowledgment cannot hope to do justice, I'd like to wish you both a very happy, forthcoming, 30th anniversary!

Donald A. Hitko
September 27, 1996

This project has been sponsored by DARPA contract DAAL-01-95-K-3526.

Table of Contents

I. Introduction	10
1.1 Recent Efforts in VCO Circuits	11
1.2 System Overview and Specifications.....	13
1.3 Scope of the Project	15
1.4 The Analog Devices RF25 (ADRF) Silicon Bipolar Process.....	17
1.5 Preliminary Measured Results	18
1.6 Simulation and Design Tools.....	21
1.7 Preview of the Thesis.....	21
II. Oscillator Theory	25
2.1 Oscillator Configurations and Considerations	25
2.1.1 Ring Oscillators	25
2.1.2 Relaxation Oscillators.....	26
2.1.3 Resonant Oscillators	27
2.1.4 Tuning the Frequency of Oscillation	27
2.2 Oscillation Criteria.....	28
2.2.1 Negative Resistance Device Configurations.....	30
2.2.2 Criteria for Oscillation	32
2.3 Resonator Circuit Models	34
2.4 Oscillator Noise Theory.....	37
2.4.1 Oscillator Noise Matching.....	39
2.4.2 Other Phase Noise Analysis Methods.....	40
III. Monolithic Inductors	43
3.1 System Considerations for the Inductors	43
3.2 Monolithic Inductor Test Wafer	44
3.3 Monolithic Inductor Wafer Test Methodology.....	46
3.4 Inductor Characterization Results.....	47
3.4.1 Detailed Data on Bond Wire Inductors.....	51
3.4.2 Lumped Element Model of S-parameter Data.....	54
3.5 Other Possibilities	56

IV. 1.8GHz (L-band) Monolithic VCO Design	58
4.1 System Considerations for the VCO.....	58
4.2 VCO Circuit Design.....	60
4.2.1 Transistor Considerations	61
4.2.2 Negative Resistance Element Topology	61
4.2.3 Noise Matching.....	65
4.2.4 Differential Oscillator Design.....	68
4.2.5 Tunable Resonator Design and VCO Interface Issues.....	75
4.2.6 The Differential Colpitts VCO Circuit	82
4.3 VCO Simulation Results.....	84
4.4 Frequency Drift.....	92
V. LMVCO Design Recapitulation and Conclusions	98
Appendix: LMVCO Chip Schematics	104

List of Figures

Figure 1-1. Simplified block diagram of frequency synthesizer.....	14
Figure 1-2. Photograph of LMVCO chip.....	16
Figure 1-3. Measured phase noise spectrum of the VCO at 1.8V, 9mW.....	19
Figure 1-4. Measured tuning range of the VCO at 1.8V, 9mW.....	20
Figure 2-1. Basic components of a negative resistance oscillator.	28
Figure 2-2. One-port negative resistance model of an oscillator.	29
Figure 2-3. DC i-V characteristics with regions of negative resistance.	30
Figure 2-4. ECP negative resistance characteristic with increasing signal power.....	32
Figure 2-5. Parallel resonator equivalent circuits using simple modeling of loss.	35
Figure 2-6. Impedance and admittance plots for a parallel resonator ($f_0=1\text{GHz}$, $Q=5$)...	36
Figure 2-7. Leeson's oscillator phase noise model.....	37
Figure 3-1. Metal mask pattern for monolithic inductor test wafer.	45
Figure 3-2. Photo of one-port bond wire inductors (1.25mm and 2mm).	46
Figure 3-3. Survey of MTL spiral inductors at 1.8GHz.	48
Figure 3-4. Survey of MTL bond wire inductors at 1.8GHz.	49
Figure 3-5. Self-resonant frequency data for MTL spiral inductors.....	51
Figure 3-6. Inductance of MTL one-port bond wire inductors at 1.8GHz.	52
Figure 3-7. Quality factor of MTL one-port bond wire inductors at 1.8GHz.	53
Figure 3-8. Lumped element model of the 2mm one-port bond wire inductor.	55
Figure 3-9. Comparison of SPICE model to s-parameter data.	55
Figure 4-1. Single-ended common-base oscillator circuit.	62
Figure 4-2. Single-ended ECP oscillator circuit.	64
Figure 4-3. Noise floor at oscillator output port with feedback ratio as a parameter.	67
Figure 4-4. Differential common-base Colpitts oscillator circuit.....	68
Figure 4-5. Effects of the feedback structure on negative resistance.	69
Figure 4-6. Emitter perimeter/area ratio effects on NPN input noise spectral density.	70
Figure 4-7. Emitter area effects on NPN input noise spectral density.....	71
Figure 4-8. Emitter area effects on NPN unity current gain frequency.	71
Figure 4-9. Emitter area effects on NPN unity power gain frequency.	73

Figure 4-10. Negative resistance of differential Colpitts with increasing signal
amplitude.....74

Figure 4-11. Impedance of differential resonator with frequency80

Figure 4-12. Schematic of differential, common-base Colpitts VCO.83

Figure 4-13. Nominal operation of VCO showing build-up of oscillations.85

Figure 4-14. Power spectrum of VCO output signal.85

Figure 4-15. Tuning range and gain of VCO.....86

Figure 4-16. Operation of VCO over a range of supply voltages with87

Figure 4-17. Operation of VCO over a range of bias currents.....89

Figure 4-18. Operation of VCO over process.....90

Figure 4-19. Calculated phase noise spectral density using Lindenmeier.....92

List of Tables

Table 1-1: VCO Targets.....	14
Table 4-1: VCO Gain Specifications	60
Table 4-2: Model Parameter Data for PNP-based Varactor Unit Cell.....	78
Table 4-3: Simulation Data for VCO Output Over Supply Voltage.....	88
Table 4-4: Simulation Data for VCO Output Over Bias Current	89
Table 4-5: Simulation Data for VCO Output Over Process	91
Table 4-6: Simulation Results of VCO Frequency over Temperature	96

I. Introduction

Boosted by the FCC's opening of the air waves to commercial, digital and spread-spectrum radio-based products, and reinforced by the mantra of recent consumer electronics spending ("Smaller, faster, better, cheaper"), the wireless communications design field is bustling with activity. Transceivers intended for RF/microwave applications such as these often call upon a voltage-controlled oscillator (VCO), a circuit which generates a signal of a frequency dictated by an input control voltage. Popular uses for such a device are as the tunable element in a frequency shifting stage, or as an integral part of a phase-locked loop (PLL). A PLL typically forms the basis for frequency synthesis and data recovery operations, and will be briefly examined later in this chapter. A widespread effort to perform these tasks more accurately while supporting current market trends has led to increased research in monolithic VCOs.

Many of the wireless applications are focused on portability, for which low power operation is a necessity if the product is to be successful. In commercial industries, success implies achieving a large customer base. For wireless data link products, having many users translates into a crowded frequency domain, thus the noise performance of the RF front-end circuitry cannot be arbitrarily sacrificed in the pursuit of decreased power and cost. The noise contributions of the elements within an oscillator lead to random fluctuations in the phase of the generated signal, resulting in a noise spectrum at the output. (Noise also affects the amplitude, but this is of a lesser concern.) When used in a wireless transmitter, the phase noise of an oscillator becomes a part of the radiated spectrum, potentially interfering with signals existing in adjacent channels—signals which may be much weaker in amplitude. For this reason, the amount of noise generated by an oscillator is often quantified by a measurement of its (phase) noise spectral density at a given offset from the carrier frequency. The offsets specified are based on system requirements and/or measurement capability, but are typically 10kHz, 100kHz, or 1MHz. However, for comparison purposes, the phase noise spectrum may be extrapolated from a single point assuming a slope of -20dB/decade. The origins of this characteristic are investigated in Chapter 2.

1.1 Recent Efforts in VCO Circuits

Much of the increased activity in this area has targeted the 902-928MHz ISM (Industrial-Scientific-Medical) band, wherein several methods of realizing an oscillator have been explored. An early work demonstrated that tunable, relaxation-based oscillators implemented in a silicon bipolar technology could reach this frequency range [1], but noise and power were not concerns in this effort. The arrival of micron and sub-micron feature size CMOS processes has carried the promise of low power and high integration levels to this band of operation, leading to the debut of 900MHz CMOS ring oscillators in 1994 [2]-[6]. Having anticipated this trend, several studies were also published over the same time frame which advanced methods for analysis of timing jitter (phase noise) in these circuits [7][8]. The best reported phase noise figure for a CMOS inverter ring voltage-controlled oscillator at 900MHz is -83dBc/Hz at a 100kHz offset, achieved with a power dissipation of 7.4mW from a 5V supply [5][6].

Despite the simplicity of this implementation, resonant oscillators have garnered the most attention as the noise levels typically required for wireless data link applications have been a limiting constraint. The most common approach to integrating a resonator has been to use passive, planar spiral inductors constructed from the metallization levels within a process. Using this technique, Duncan, et.al., have developed a 5V oscillator with quadrature outputs in the 1GHz range, which consumes about 80mW in a bipolar circuit to attain a -87dBc/Hz phase noise level at an offset of 100kHz [9]. Ali and Tham also used spiral inductors with an emitter-coupled pair active element to operate a 3V, 10mW oscillator at 900MHz which yielded -101dBc/Hz at 100kHz from the carrier [10].

Realizing that most of the resonator loss occurs in the substrate as RF energy couples into it from the inductor, several groups have used spiral inductors from beneath which the silicon has been etched. A back-side etch performed the selective substrate removal for the 1GHz oscillator of Basedau and Huang, which was implemented in a 1 μ m CMOS process [11]. A single-ended output with a noise spectral density of -95dBc/Hz (100kHz offset) was produced with 1.5V supplying 16.5mW. By a different processing technique—using a top-side gaseous etch to clear cavities beneath the spiral inductors—Rofougavan, et. al, designed a quadrature oscillator, also with CMOS gain elements (1 μ m NMOS transistors) [12]. Consuming 30mW from a 3V supply, this VCO yielded a noise

spectrum 85dB down from the carrier (1 Hz integration bandwidth), at an offset 100kHz away from a 900MHz center frequency.

A similar path has evolved among the efforts in the multi-GHz range (e.g., 1.8GHz and 2.4GHz applications). Soyuer and Warnock pushed relaxation oscillators up to 5GHz in a 0.8 μ m bipolar process [13], but did not characterize the phase noise performance of their circuit which dissipated 70mW. In an attempt to obtain better performance while maintaining a tightly integrated circuit, Aytur and Razavi developed a ring oscillator composed of two differential gain stages with bipolar transistors [14]. Low power levels were reported (3mW in the VCO from a 3V supply), but the phase noise characteristics were comparatively poor, with measured spectral densities of -74dBc/Hz and -88dBc/Hz at spot frequencies 1MHz and 4MHz from the carrier, respectively.

Perhaps even more so in this higher frequency range, inductorless designs have remained a novelty, while spiral inductors in lumped element resonators have proven to be the primary means of implementation. Perhaps the original work in monolithic, 1.8GHz, harmonic VCOs was by Nguyen and Meyer [15], an unusual design in that a tuning element (e.g., a varactor) was not involved, but rather the control voltage was used to steer current between two LC resonators with center frequencies split 180MHz apart. The oscillation was smoothly tunable between the two center frequencies, and achieved 100kHz offset phase noise levels of -88dBc/Hz while dissipating 70mW at 5V. Soyuer, et. al., have taken spiral inductors to extremes, first with a four-level metal bipolar implementation of a 2.4GHz oscillator [16], then with a 4GHz VCO using a CMOS gain stage in a five-level metal process [17]. The bipolar circuit generated a single-ended output using 50mW from a 3.6V supply, creating -92dBc/Hz of noise 100kHz away from the 2.4GHz carrier. The CMOS VCO used a source-coupled pair in a 3V, 18mW circuit, with the 4GHz output being developed across a 50 Ω resistor in one drain terminal, and the resonator located in the other. Noise performance was measured to be -106dBc/Hz at an offset of 1MHz.

An alternative approach to on-chip resonators was advanced by Craninckx and Steyaert, who used pad-to-pad bond wires as inductors [18]. With this technique, a VCO built in a CMOS process reached 1.8GHz, having a phase noise level -85dBc/Hz just 10kHz from the carrier. Operating with a supply of 3V, this circuit dissipated 24mW. As

a point of comparison, Piazza and Huang describe a 1.6GHz VCO using an external resonator for GPS (Global Positioning System satellite) receivers [19]. While the nature of the digital spread-spectrum signaling in this application requires only modest noise performance, the rationale offered for this strategy was that a savings in power could be realized with the higher Q, off-chip components. With a single-ended circuit implemented in a bipolar process, measured noise performance was -95dBc/Hz, 100kHz away from the oscillation, at an energy consumption rate of 3.2mW from a 3V supply.

1.2 System Overview and Specifications

The VCO being designed as the core of this thesis is an integral part of a PLL-based frequency synthesizer which, along with a power amplifier to drive the antenna and a bandpass filter to remove harmonics induced by amplifier non-linearity, forms the high data rate RF transmitter for a low power, wireless sensors research program at MIT. A simplified block diagram of the transmit function is shown on the following page in Figure 1-1. The VCO produces an output signal Φ_o , of frequency $\omega_o = \frac{d\Phi_o}{dt}$, which is divided down (in terms of frequency), and then compared with a fixed frequency reference (Φ_{REF}) source. The output of the phase detector—fundamentally the result of a logical XOR operation—is a time periodic signal, the DC component of which is proportional to the difference between the phase of the reference and the phase of the scaled VCO output. The lowpass filter then removes the high frequency components (e.g., $\frac{d}{dt}(\Phi_{REF})$) from the phase detector output, leaving the low frequency signal V_c , the control voltage for the VCO— $V_c = K_p\left(\Phi_{REF} - \frac{\Phi_o}{N}\right)$, where K_p is the gain of the phase comparator in Volts per radian. The VCO output signal can then be expressed as $\Phi_o(t) = K_v \int V_c(t) dt$, where K_v is the oscillator gain in radians per second per Volt. For further discussion of the basics of phase-locked loops, the components within them, and the design of frequency synthesizers from them, the reader is referred to any number of modern texts such as Crawford [20].

When the PLL is in lock, the frequency of the output signal will be a multiple of the reference frequency, the multiplication factor being the localized time-average value of the divider. In the proposed high rate transmitter [21], the digital data modulates the RF signal by dynamically selecting the divide value via the modulus controller. The workings of this controller and the implementation details of the frequency synthesizer are well

beyond the scope of this document. Of primary importance for the VCO, however, are the frequency (ω_o) range over which it will oscillate, the linearity of its transfer characteristic (i.e. how constant K_v remains over the tuning range), and the spectral purity of the VCO output (i.e. phase noise).




Figure 1-1. Simplified block diagram of frequency synthesizer.

Table 1-1: VCO Targets

CHARACTERISTIC	VALUE
Frequency band of oscillation	1.8GHz
Phase noise spectral density, 1MHz offset	-140dBc/Hz
Phase noise spectral density, 100kHz offset	-120dBc/Hz
Power dissipation	< 20mW

Exact specifications for the VCO have not been penned; it is the RF transmitter which must be optimized for the sensor application, not the oscillator by itself. Nonetheless, in conjunction with simulations of the frequency synthesizer, approximate targets for the VCO have evolved, and are shown above in Table 1-1. System requirements call for the capability of handling data at a rate that may vary continuously from 1 bit/s to 1 Mbit/s, and the figures listed are based on the maximum data rate. The phase noise density at a 1 MHz offset was calculated by providing for an 80 dB signal-to-noise ratio in the output signal; the noise at an offset of 100 kHz is provided for reference purposes, and is derived by allowing for the typical -20 dB/decade slope of the output power spectrum. These

numbers are optimistic, being more stringent than current standards such as GSM and DECT, but have been crafted with the intention of pushing technology. Definitions will become more precise as the wireless sensor system matures and unfolds so that performance-optimizing trade-offs can be made between system blocks.

1.3 Scope of the Project

As a piece of MIT's wireless sensor project, a collaborative research effort forged to further low-power circuit design techniques for wireless applications, the goal of this thesis is to develop a monolithic VCO which operates around 1.8GHz, is compatible with other required electronics, and has power consumption and phase noise characteristics that are at the leading edge of the art. Efforts entail the design, simulation, and layout of an integrated circuit in a silicon bipolar process—hereafter referred to as the L-band Monolithic Voltage-Controlled Oscillator (LMVCO) chip—containing the oscillator and other circuitry needed to fulfill system and testing requirements. A die photo shown on the following page in Figure 1-2 illustrates the work, in which the VCO has been identified along with input circuitry, a buffer/frequency divider (divide by 2), 50 Ω output driver for a balanced 900MHz signal, and the inductors used in the oscillator. The buffer/divider circuit drives a pair of coplanar waveguide probe pads from which the differential 1.8GHz signal can be sensed. Also included on the finished silicon are probable test cells with the capacitors, varactors, and transistors singled out from the VCO. (The inductors are handled on other die sites.) Four LMVCO chip versions evolved over the course of this project to explore new varactor devices and tuning mechanisms; these are explained in Chapter 4.

Eventually, a single chip implementation of the transmitter will be sought, but in order to expedite development during the initial phases of the project, research has been focused on smaller units. The remainder of the phase-locked loop which controls the VCO has been constructed in a 0.6 μm CMOS technology, in which handling a signal transitioning at 1.8GHz has proven to be difficult [21]. To overcome this limitation, the first stage of the prescaler is accomplished on the LMVCO chip (the divide by 2 circuit). This effectively places the RF output of the frequency synthesizer at the on-chip ground-signal-ground RF probe pads, while the feedback chain of the PLL is completed via the 900MHz signal produced by the bipolar frequency divider stage. This reduced speed signal is then

driven off the LMVCO chip by the 50Ω output amplifier, onto a microstrip trace feeding the CMOS transmitter IC that houses the rest of the prescaler, phase/frequency detector, and lowpass filter. The loop is then closed by communication in the reverse direction of the baseband data stream, serving as the control input to the VCO.

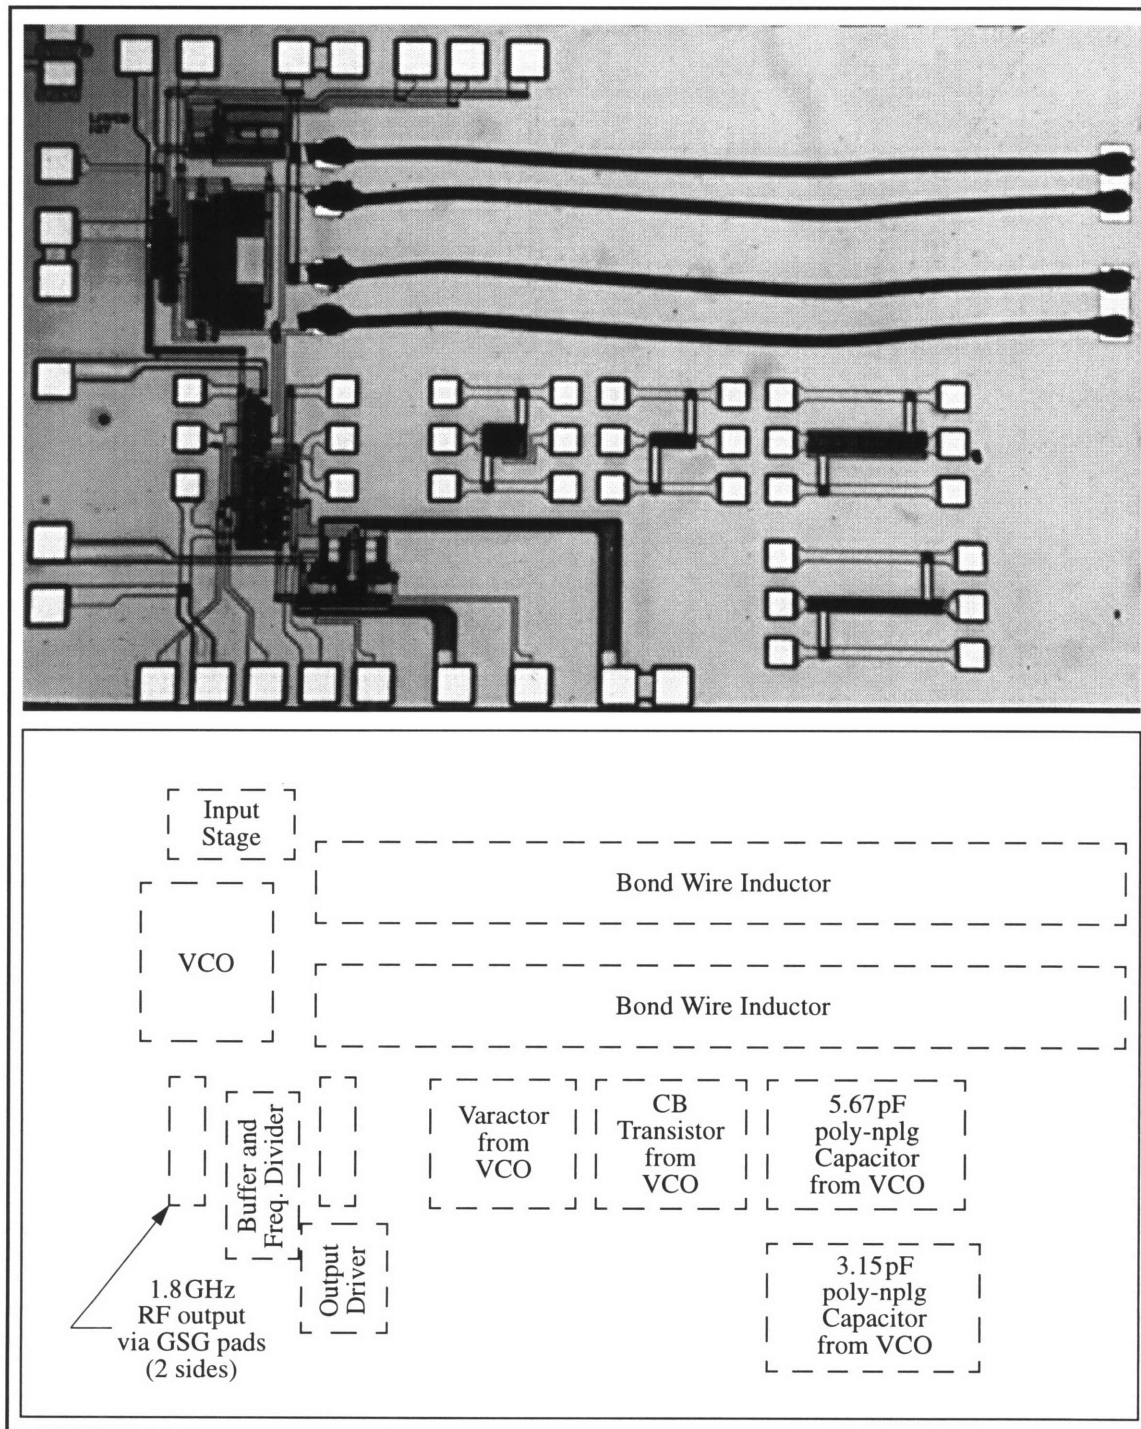


Figure 1-2. Photograph of LMVCO chip.

1.4 The Analog Devices RF25 (ADRF) Silicon Bipolar Process

Achieving the wireless sensor project goals of small size and low power operation points to the need to realize a system with high levels of integration. Although microwave integrated circuits have been the nearly exclusive province of GaAs technologies over the previous decade, continual improvements in silicon wafer processing have pushed the performance envelope of silicon-based devices to the point where they are now a feasible alternative for many applications in the low GHz range. With increasing commercialism bringing cost, size, and power conscious market opportunities into these frequency bands, silicon technologies are rapidly gaining a foothold. Whenever these issues are present, silicon inevitably becomes the choice provided circuit and system performance requirements can be met.

The power dissipated in RF circuitry is generally of the static variety, either in the DC drive levels needed to support device and circuit operation at the required frequencies, or in converting to RF energy. This is the case with most microwave oscillators, typified by those presented in Section 1.1; thus there is no inherent power savings to choosing one transistor type over the other (CMOS or bipolar). For VCOs then, the decision will primarily be based upon the needs to integrate other forms of circuitry (e.g., power amplifiers, digital signal processing), although bipolar transistors in current processes generally yield higher operating frequencies for a given bias level. Restated in terms more directly relevant to the MIT wireless sensors project, this suggests that BJTs should allow for a lower power consumption in meeting system requirements. Similar conclusions can be drawn regarding noise; compared with CMOS, bipolar devices generally have both lower noise figures at the frequencies of interest, and also reduced levels of surface and recombination effects which give rise to flicker noise.

With these factors to consider, Analog Devices' RF25 (ADRF) technology is a good choice for this project. The cornerstone of this process is a double-polysilicon, self-aligned, NPN transistor suitable for 5V operation with a peak f_T of 25GHz, reachable with a current density of $0.3\text{mA}/\mu\text{m}^2$. To reduce base resistance—a major contributor to noise figure—p+ base polysilicon is used to make contact to a heavily-doped extrinsic base region in the single-crystal silicon that is deeper than the intrinsic base width. This extrinsic diffusion pocket surrounds the active device area to minimize base (link) resistance.

Yet another limiting device factor in microwave circuit design is the collector-to-substrate capacitance of the transistor, the presence of which degrades high frequency performance. Forming the collectors in a lightly doped epitaxial layer ($10\Omega\text{-cm}$, p-type) on top of an even more lightly doped p-type wafer ($20\Omega\text{-cm}$) helps to minimize this capacitance, as does generous spacing between the collector structure and the p-isolation regions which encircle the devices.

Conventional lateral PNP transistors are offered in ADRF, which use the NPN base handle for collector and emitter contact, while the NPN collector region, n+ buried layer, and n+ collector tie (n+ plug) serve as the PNP base terminal. The resulting device works well enough to obtain peak values for f_T and the βV_A product of 70MHz and 1000, respectively. To complete the active device arsenal, NMOS transistors with $0.7\mu\text{m}$ effective channel lengths are also available. An unsilicided, p+ layer is used for the gate material, resulting in higher than typical gate resistances and threshold voltages, but the MOSFET is well suited for RF switching duty.

The same p+ layer used for the n-channel gates is also used in several passive devices. The capacitors within the ADRF stable use this layer as the top plate over 230\AA of oxide, grown on a bottom plate made from the n+ collector plug implant. A capacitance per unit area of $1.5\text{fF}/\mu\text{m}^2$ is yielded by this MOS structure, with a ratio of 15:1 between this value and the bottom plate parasitic. Two p+ polysilicon resistors are also supported. To realize lower resistances, the upper plate of the capacitor provides $150\Omega/\square$ with a temperature coefficient of $300\text{ppm}/^\circ\text{C}$. A compensation implant is available to increase the resistance to $800\Omega/\square$, but with wider process and temperature variation ($-800\text{ppm}/^\circ\text{C}$). Further process information and characterization data have been documented in an article published by Analog Devices [22].

1.5 Preliminary Measured Results

The LMVCO chips have been designed and fabricated in this process. As of this writing, wafers have been received, and some preliminary measurements have been conducted. Diced VCOs are placed in open-lid RF packages, with non-critical DC signals wire bonded to package pins; remaining signals are probed on chip. Cascade Air Coplanar™ microwave probes are used to pick up the 1.8GHz signal. Designed to work with a

supply of 1.8V, 5mA (including a bias stage), the VCO has proven functional at this 9mW level, achieving a measured noise density of -118dBc/Hz at a 1MHz offset. The remainder of the SSB phase noise spectrum is shown in Figure 1-3. This measurement was taken open loop with an HP8563E spectrum analyzer equipped with the phase noise utility, which has a noise floor of about -133dBm. The noise floor is reached by the VCO spectrum at an offset of 3MHz from its -10dBm carrier. Glitching in the measurement at 100kHz was later traced to a piece of laboratory equipment that was located adjacent to the spectrum analyzer, and is not inherent in the LMVCO chip. The continuity of the trend to either side of the 100kHz offset also supports this assertion. Finally, the close-in characteristic is a result of some modulation occurring in the open loop VCO. Closed loop measurements may be taken to exorcise this source of error from the data.




Figure 1-3. Measured phase noise spectrum of the VCO at 1.8V, 9mW.

While this performance is no small accomplishment, the series of LMVCO chips was designed as a low power exploration vehicle, so it is natural to inquire as to the extent

to which consumption may be reduced. Lowering the supply to 1.5V demonstrated no appreciable decline in performance, although the noise spectrum began to rise for values smaller than this. Similarly, degradations in oscillation amplitude and noise performance were observed for reductions in bias current. With either a 1.5V or 1.8V supply, operation was sustained with 3.1 mA of VCO current, but with phase noise readings hovering around -100dBc/Hz, 1MHz away from the carrier. Oscillation is supported with power levels as low as 3.2mW with a supply of 0.9V, yielding -96dBc/Hz at the same 1MHz offset.

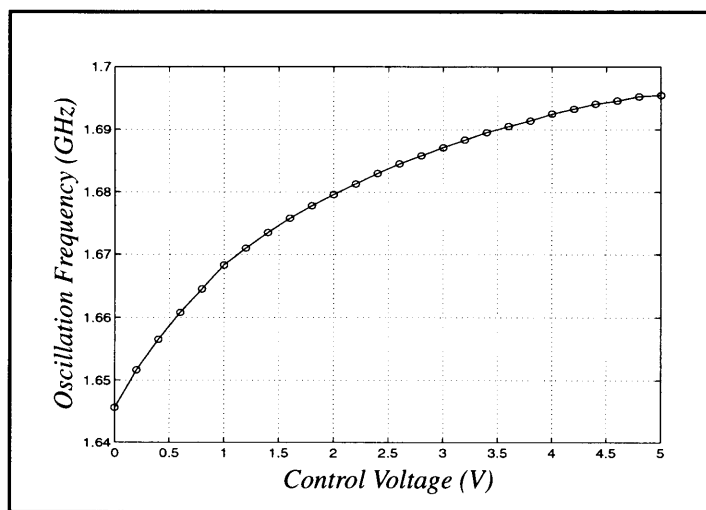


Figure 1-4. Measured tuning range of the VCO at 1.8V, 9mW.

The tuning range data for the VCO is graphed above in Figure 1-4. As used in the frequency synthesizer, the input signal to the oscillator will be bounded by the 3V supply of the operational amplifier which drives it, but here the oscillation frequency has been gauged out to 5V to demonstrate the capability of ADRF to handle larger potentials. At the middle of the input control voltage range, the VCO gain is 12MHz/V. This is somewhat below the nominal value of 20MHz/V intended for this VCO design, a fact which is likely correlated with the low center frequency observed in preliminary parts. As previously mentioned, this oscillator is designed to operate at 1.8GHz, but instead has been seen to tune only from 1.646GHz to 1.696GHz. While this variation is within process tolerances (as demonstrated in Chapter 4), the device test structures on LMVCO chip silicon have not yet been probed, so it remains unclear whether the lower oscillation frequency is

due to something systematic in the design, or merely owed to capacitance values toward the upper end of their statistical ranges. Further laboratory evaluation will be conducted to determine the source of this error, and to better characterize the low power operation of the oscillator. Although not a priority for the first turn of silicon, tuning mechanisms to allow for correction of process variations will be considered for future revisions of the VCO.

1.6 Simulation and Design Tools

The primary aid for the circuit design work in this project has been ADICE (Analog Devices Integrated Circuit Emulator). SPICE-like linear and transient analyses are supported with an extended Gummel-Poon model of bipolar transistors. Enhancements for simulating thermal, weak avalanche, quasi-saturation, and Early voltage bias dependence effects have been implemented, and the models have been expanded for more accurate handling of parasitic elements in the epitaxial collector region. Modeling data have been encoded into scalable “super-models” written for ADICE, which calculate model parameters for custom layout devices. Using this approach, the ADICE netlister enters a unique subcircuit call for each element in the schematic (captured with Cadence’s Composer), which then access parameterized .SUBCKT definitions in ADICE libraries. These definitions include a network of devices to account for parasitic effects in the NPN transistors and passive circuit elements.

While this helps to ensure accurate circuit simulation, it is often not the quickest, nor most theoretically conducive, design tool. Circuit dependencies and optimization methodologies have been developed and explored with linear circuit models using the general-purpose MATLAB® and Maple® mathematics programs. Specific examples in this work include the analysis of an optimum noise match to the gain stage in the oscillator, and the effects of changes in the positive feedback structure upon the negative resistance element; both of these techniques are examined in Chapter 4.

1.7 Preview of the Thesis

Chapter 4, the last in this document, details the circuit design of the LMVCO chip. The impact of system considerations on the design of the VCO are discussed, from

which the choices of circuit devices and topology are unfolded. A linear noise matching analysis is presented, along with the simulation design data used to pen the circuit parameters. The tunable resonator is tackled next, including versions with experimental varactor structures. This chapter concludes with extensive simulation results, along with a brief treatise on frequency drift.

The theoretical foundation for this design work is built in Chapter 2. Oscillator configurations are pondered, and the criteria for stable oscillation are presented for the harmonic class of oscillators (i.e. those using a resonant element). An overview of resonator circuit models is provided as a prelude to a discussion of the existing theories of oscillator noise. This serves to illustrate the importance of high quality passive components, of which on-chip inductors have generally been a key performance barrier. An empirical look at methods of improving integrated inductors is the topic of Chapter 3.

Although each of the remaining three chapters covers a vital part of the oscillator design—from the theoretical, to the pragmatic, and finally the circuits and simulations themselves—each is constructed to stand on its own. Cross-references are provided where useful, however, an attempt was made to carry along sufficient introductory and explanatory material so the reader may delve into any aspect of the work without a lot of backtracking. Throughout this thesis, each unit is internally summarized, and reference lists are provided at the end of each chapter. Hopefully, this approach will make the information contained within to be more easily accessible.

References

1. J. Wu, “A Bipolar 1 GHz Multi-Decade Monolithic Variable Frequency Oscillator,” *ISSCC Digest of Technical Papers*, International Solid-State Circuits Conference, Feb. 1990, pp. 106-107.
2. J. Min, et. al., “An All-CMOS Architecture for a Low-Power Frequency-Hopped 900MHz Spread Spectrum Transceiver,” *IEEE Custom Integrated Circuits Conference Proceedings*, 1994, pp. 16.1.1-16.1.4.
3. M. Thamsirianunt and T. Kwasniewski, “A 1.2 μ m CMOS Implementation of a Low-Power 900MHz Mobile Radio Frequency Synthesizer,” *IEEE Custom Integrated Circuits Conference Proceedings*, 1994, pp. 16.2.1-16.2.4.

4. H. Liu, et. al., "A Low-Power Baseband Receiver IC for Frequency-Hopped Spread Spectrum Applications," *IEEE Custom Integrated Circuits Conference Proceedings*, 1995, pp. 14.2.1-14.2.4.
5. T. Kwasniewski, et. al., "Inductorless oscillator design for personal communications devices—A 1.2 μ m CMOS process case study," *IEEE Custom Integrated Circuits Conference Proceedings*, 1995, pp. 14.6.1-14.6.4.
6. M. Thamsirianunt and T. Kwasniewski, "CMOS VCOs for PLL Frequency Synthesis in GHz Digital Mobile Radio Communications," *IEEE Custom Integrated Circuits Conference Proceedings*, 1995, pp. 14.7.1-14.7.4.
7. B. Kim, et. al., "PLL/DLL System Noise Analysis for Low Jitter Clock Synthesizer Design," *Proceedings of the IEEE International Symposium on Circuits and Systems*, June 1994, pp. 31-34.
8. T. Weigandt, et. al., "Analysis of Timing Jitter in CMOS Ring Oscillators," *Proceedings of the IEEE International Symposium on Circuits and Systems*, June 1994, pp. 27-30.
9. R. Duncan, et. al., "A 1 GHz Sinusoidal Oscillator," *IEEE Custom Integrated Circuits Conference Proceedings*, 1995, pp. 91-94.
10. A. Ali and J. Tham, "A 900MHz Frequency Synthesizer with Integrated LC Voltage-Controlled Oscillator," *ISSCC Digest of Technical Papers*, International Solid-State Circuits Conference, Feb. 1996, pp. 390-391.
11. P. Basedau and Q. Huang, "A 1GHz, 1.5V Monolithic LC Oscillator in 1- μ m CMOS," *European Solid-State Circuits Conference Proceedings*, 1994, pp. 172-175.
12. A. Rofougaran, et. al., "A 900MHz CMOS LC-Oscillator with Quadrature Outputs," *ISSCC Digest of Technical Papers*, International Solid-State Circuits Conference, Feb. 1996, pp. 392-393.
13. M. Soyuer and J. Warnock, "Multigigahertz Voltage-Controlled Oscillators in Advanced Silicon Bipolar Technology," *IEEE Journal of Solid-State Circuits*, April 1992, pp. 668-670.
14. T. Aytur and B. Razavi, "A 2-GHz, 6-mW BiCMOS Frequency Synthesizer," *IEEE Journal of Solid-State Circuits*, December 1995, pp. 1457-1462.
15. N. M. Nguyen and R. G. Meyer, "A 1.8 GHz Monolithic LC Voltage-Controlled Oscillator," *IEEE Journal of Solid-State Circuits*, March 1992, pp. 444-450.
16. M. Soyuer, et. al., "A 2.4-GHz Silicon Bipolar Oscillator with Integrated Resonator," *IEEE Journal of Solid-State Circuits*, February 1996, pp. 268-270.

17. M. Soyuer, et. al., "A 3 V 4GHz nMOS Voltage-Controlled Oscillator with Integrated Resonator," *ISSCC Digest of Technical Papers*, International Solid-State Circuits Conference, Feb. 1996, pp. 394-395.
18. J. Craninckx and M. Steyaert, "A CMOS 1.8 GHz Low-Phase-Noise Voltage-Controlled Oscillator with Prescaler," *ISSCC Digest of Technical Papers*, International Solid-State Circuits Conference, Feb. 1995, pp. 266-267.
19. F. Piazza and Q. Huang, "A 12mA Triple-Conversion Receiver for GPS," *ISSCC Digest of Technical Papers*, International Solid-State Circuits Conference, Feb. 1996, pp. 286-287.
20. J. Crawford, *Frequency Synthesizer Design Handbook*, Artech House, Norwood, MA, 1994, Chapters 2 and 4.
21. M. Perrott, "Techniques for High Data Rate Modulation and Low Power Operation of a Fractional-N Frequency Synthesizer with Noise Shaping," Proposal for Thesis Research, Massachusetts Institute of Technology, April 1996.
22. K. O, et. al., "A Low Cost and Low Power Silicon *n*pn Bipolar Process with NMOS Transistors (ADRF) for RF and Microwave Applications," *IEEE Transactions on Electron Devices*, October 1995, pp. 1831-1840.

II. Oscillator Theory

A fundamental element of most RF transceiver architectures is a voltage-controlled oscillator (VCO), an element which generally consumes a significant portion of the overall power budget in low power applications. In addition, the performance of the VCO, measured in terms of its output noise spectrum, is a determining factor in several key limitations and capabilities of modern wireless data transmission systems. Although they are used in many applications, oscillators—particularly those used in communications circuits—often lie hidden behind a symbol in a block diagram, emblazoned in an aura of mystique. Due to their non-linear nature, oscillators tend to defy the normal methods of circuit analysis, leaving in their wake almost as many models and design techniques as there are circuit topologies. The diversity which has resulted should aid the designer; just as each oscillator configuration works better in some applications than others, certain models fit some circuits better than others. Choosing among these possibilities constitutes a large portion of the “art” of oscillator design.

2.1 Oscillator Configurations and Considerations

At the broadest level, oscillators can be thought of as existing within two classes: those based upon the controlled relaxation of an energy storage device, and those regulated by the frequency response of a resonant element. The ring oscillator, a special type of relaxation oscillator using only parasitic storage elements, deserves notice unto itself due to the unique application domains and design methodologies which apply to it. Some forms of each of these three configurations may be integrated with a variety of processes, devices, and techniques, providing a well-rounded arsenal to the circuit designer.

2.1.1 Ring Oscillators

A ring oscillator is constructed from a string of inverting stages, with the output of the chain tied back around to the input in an astable fashion (i.e. an odd number of inversions exists within the string). Enough delay is built into the chain such that—in the absence of severe loading—each node completely transitions between low and high states.

Since the only energy storage devices within the ring are the parasitic elements of the transistors themselves (plus the wiring), this circuit freely runs at a speed limited only by the gate delay of a process and the number of stages needed to achieve the required delay in the chain of at least 180° . This astable mode of operation generally calls for at least three stages in the ring (more can be used to operate at slower frequencies), but two inverting stages have been used [1] with a differential circuit topology.

CMOS ring oscillators, in particular, find common use alongside digital circuitry in moderate frequency applications. However, despite recent processing advances, the desired 1.8GHz band remains a difficult target. This is a result of the nature of device scaling, which calls for a reduction in supply voltage to accompany the scaling of channel lengths to maintain reliability. As a result, overall gate transition times have not been scaling at the same rate. Even if this frequency were reachable, recent research has indicated that the CMOS ring oscillator approach would not meet the phase noise and power targets for this application. Although it has been shown that timing jitter should be reduced with shorter channel lengths (due to a reduction in gain) [2], work with CMOS ring oscillators at 900MHz [3] indicates that the phase noise remains high compared with other configurations. Another key limitation of the CMOS inverter approach is that power consumption is proportional to frequency; thus it is not clear that this technique can yield low power, high frequency oscillators.

2.1.2 Relaxation Oscillators

Instead of relying on node parasitics to slew circuit switching, an energy storage element is explicitly added to the circuit in other forms of relaxation oscillators. The charging and discharging of this device, typically a capacitor, by controlled currents generates a well-determined voltage characteristic which may be used as a timing base. These multivibrator circuits are easily integrated, and find widespread use in low to moderate frequency applications such as the classic 555 IC timer. Recent efforts have extended emitter-coupled multivibrators into the GHz range [4][5], but good noise performance and the fast slewing rates that are required come only at the expense of high power consumption. Another potential downside of using relaxation oscillators for wireless purposes is that they generate waveforms with high harmonic content, a side effect which may necessitate increased filtering in both transmit and receive signal paths.

2.1.3 Resonant Oscillators

Circuits which produce a sinusoidal (or nearly sinusoidal) excitation as determined by a narrowband filter—a resonator—are termed resonant, or harmonic, oscillators. The resonant device may consist of a crystal to provide for oscillation in the low RF range, lumped inductive and capacitive components for operation up to a few GHz, or distributed elements and cavities for use in the microwave and millimeter wave bands. Regardless of the frequency range, the chosen resonator needs to be coupled to an active element which supplies energy to support the oscillations. Many devices and topologies are possible, of which some of the more commonly used transistor based configurations have been given names such as Hartley, Pierce, Clapp, and Colpitts. These names refer to distinctions in the implementation of positive feedback around the gain stage and how the resonator is coupled to the resulting arrangement; however they remain quite general terms which do not imply the use of any particular circuit.

As an integral part of the frequency synthesizer for a wireless transmitter, a VCO is required which simultaneously operates at low power levels and minimizes the output spectrum due to phase noise while being monolithically compatible with other required electronics. Resonator based oscillator architectures offer the best performance at microwave frequencies along with minimal active element requirements; the challenge, however, remains finding a monolithic resonator of sufficient quality. This issue is addressed in part by the subsequent chapter dealing with monolithic inductors.

2.1.4 Tuning the Frequency of Oscillation

To this point, the discussion has not made mention of the need to vary the frequency of oscillation. This operation is straightforward with relaxation (including ring) oscillators, where the most popular method is to modulate the current which charges/discharges the storage element. A voltage-to-current converter can be prepended to achieve a voltage-controlled oscillator. Some relaxation architectures rely upon comparators or Schmitt triggers to drive the toggling between states; these circuits allow an additional possibility of changing the frequency by manipulating the (voltage) trip point at which the transition takes place.

The situation is somewhat more involved for resonant oscillators wherein a change in frequency calls for alteration of the value of a passive component. Instruments exist

which allow this to be done mechanically, but active, electronic control is required for many applications. Either a tunable inductor or a variable capacitor is needed, and the latter option—termed a varactor—is the most oft-used. This device exploits the voltage-dependence of a p-n junction capacitance under reverse bias. To the first-order, the voltage-capacitance characteristic goes as $C \propto V^{-m}$, where $m \cong 1/3$ for linearly graded devices, but approaches $m = 1/2$ for abrupt junction doping profiles. This varactor relationship generally results in a narrow tuning range and a highly non-linear frequency versus voltage gain curve. Active inductors have been tried in an attempt to improve upon these limitations [6][7], but often have frequency constraints and have consumed significant power. As a result, most modern narrowband systems make use of a varactor (e.g., [8][9][10]), wherein efforts are then placed into designing circuits and specifications around the V^{-m} characteristic.

2.2 Oscillation Criteria

At higher RF and microwave frequencies—where only the simplest oscillator topologies predominate—the negative resistance port model is generally the least cumbersome to use, and carries the additional benefit of being easily compatible with s-parameter device characterization. For these reasons, emphasis here is upon the negative resistance viewpoint, although feedback models may be equivalently applied to much of the discussion. Feedback methods are developed in many texts, such as Clarke and Hess [11].

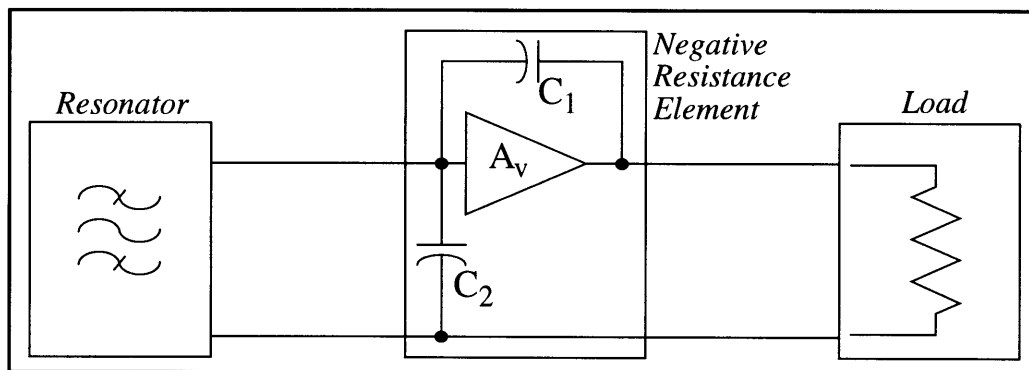


Figure 2-1. Basic components of a negative resistance oscillator.

A two-port model of an oscillator is shown in the diagram above, and contains three basic components: a frequency-selective device (resonator), an active element which

creates poles in the right-half of the s-plane (or, equivalently, a negative resistance), and some kind of loading or termination. In the example of Figure 2-1, the negative resistance element is constructed by placing positive feedback around a gain stage, such that the impedance looking into one of the ports has a negative real part at the frequencies of interest. To this port, the resonator is connected, which largely determines the frequency of oscillation. Termination of the other port must be chosen to ensure the existence of the negative resistance region. The output may be taken from either port, since it may be shown [12] that if one port is oscillating, so too is the other. Loading the input (negative resistance) port effectively lowers the quality factor of the resonator, and can also lead to frequency pulling. Placing the load at the output (termination) port may instigate the need for a matching network, and may make guaranteeing oscillation difficult if the output power or loading varies.




Figure 2-2. One-port negative resistance model of an oscillator.

One-port negative resistance elements may also be used as shown in Figure 2-2 (in this example, the diode could be a Gunn diode or an IMPATT diode—refer to Sze [13] or Carr [14] for these and other options). The placement of the load should match the configuration of the resonator and the negative resistance element, an important detail which will be discussed further. By folding the termination of a two-port device into the active element, two-port oscillators may be collapsed down to this simplified model. Here, the admittance looking into the active device is denoted $Y_a(s)$, and that in the opposite direction is labeled $Y_r(s)$.

2.2.1 Negative Resistance Device Configurations

The above discussion, and figures which accompany it, used admittances— $Y_a(s)$ and $Y_r(s)$ —to characterize the circuit. This selection is not arbitrary. In the handling of this topic, many references—probably originating with Kurokawa [15]—have been made to distinctions between “voltage-controlled negative resistance” elements (VCNRs) and those which are “current-controlled” (ICNRs). Unfortunately, little clarification has generally accompanied these references, which are most likely derived from DC i-V characteristics. A VCNR would describe a device in which the current through it was a single-valued function of the voltage across it. The situation is reversed for ICNRs; an indicative plot of each is shown in Figure 2-3. These gauges, however, are of little merit when the active device exhibits no negative resistance at DC. Nor do these definitions help to convey the importance of the distinctions.

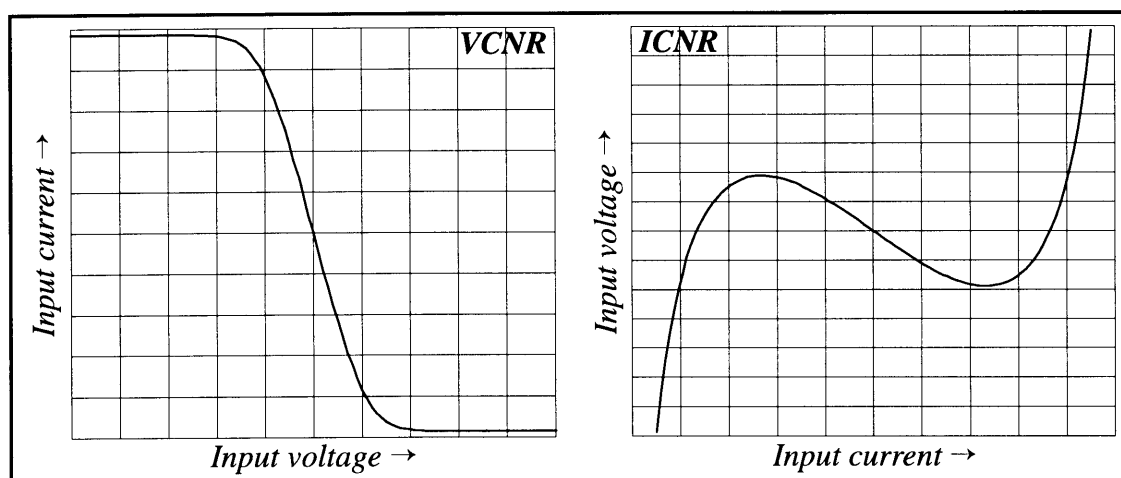


Figure 2-3. DC i-V characteristics with regions of negative resistance.

A more coherent approach may be embarked upon by considering a fundamental assumption of resonant oscillators—that it is non-linearities (i.e. limiting) on the part of the active element which set the amplitude of the oscillation, an occurrence without which the oscillation would grow without bound. The significance is held by the manner in which the amplitude of the signal is limited—whether the negative resistance element clamps the current levels or the voltage levels of the signals impressed upon it. Therein lies the assumption made by Kurokawa [15] in his landmark treatise on negative-resistance oscillators: that either the voltage across the active element is “near-sinusoidal”

(i.e. the current is clamped while the voltage swings freely—a VCNR), or else the current through the device is unhindered (ICNR), and the voltage is clamped instead. From each of these assumptions are drawn parallel sets of oscillation criteria, which are presented in the section which follows.

At this point, however, the designer may be left flagging, still without a solid means by which to decipher between these configurations, and therefore unable to discern which oscillation criteria are to be fulfilled. In some circuits, the distinction may be readily apparent, such as with an emitter-coupled pair with positive feedback (e.g., see Ali and Tham [8]), wherein the current is limited to that supplied by the tail source. When the feedback is DC (the positive input and output are hard-wired to form a one-port negative resistance element), the i - V curve which results is the VCNR shown previously in Figure 2-3. Even in such circuits, however, objective proof should be sought. One good methodology for achieving numerical verification has been suggested by Niehenke [16]. A signal source of the desired oscillation frequency is placed to drive the negative resistance port of the active element, and then both the impedance and admittance of the device are charted as a function of increasing signal amplitude. A situation where the magnitude of the real part of the impedance ($|R|$) at this frequency is seen to decrease monotonically with amplitude indicates that the voltage is being limited. This case is consistent with the ICNR definition. Reciprocally, a monotonic decrease in the magnitude of the conductance is consistent with a VCNR. For the emitter-coupled pair element discussed earlier, the impedance and admittance as a function of signal power were calculated with Libra and are plotted in Figure 2-4. As can be seen, this circuit is characterized as having a negative conductance which decreases with signal power, while the value of the negative resistance increases over the same range.

Additional cautionary notes should be appended, however, since if the magnitude of the reactance (of the active device) is large, *both* $|R|$ and $|G|$ can decrease. It may be shown that given a constant (or nearly constant) value of X , $|R|$ and $|G|$ are directly proportional (i.e. both increase or both decrease) for the condition $|R| < |X|$. The key here is monotonicity; the significance of this, as well as the VCNR/ICNR distinction, will become more translucent in the next section.




Figure 2-4. ECP negative resistance characteristic with increasing signal power.

2.2.2 Criteria for Oscillation

Provided that the harmonic content is low in the voltage across the active element, Kurokawa [15] demonstrated that in the condition of steady-state oscillation, the complex conjugate poles are on the $j\omega$ -axis at the frequency where the imaginary parts of $Y_a(s)$ and $Y_r(s)$ sum to zero. Mathematically, the condition of stable oscillation applicable to a voltage-controlled negative resistance may be stated as:

$$\Re \{ Y_a(j\omega_0) \} + \Re \{ Y_r(j\omega_0) \} = 0$$

$$\Im \{ Y_a(j\omega_0) \} + \Im \{ Y_r(j\omega_0) \} = 0,$$

which then may be solved for the frequency. (For simplicity, the amplitude dependence of the active element has been dropped.) To ensure that the proper oscillations build from initial conditions, it is necessary that the poles originate in the right-half of the s-plane so that the circuit is unstable about the initial operating point. Including this requirement, and recasting the equations in terms of conductance and susceptance yields:

$$G_a(\omega_0) + G_r(\omega_0) < 0$$

$$B_a(\omega_0) + B_r(\omega_0) = 0.$$

Kurokawa also illustrated that once the desired condition of steady-state oscillation is reached, the circuit must then be stable about the resulting operating point. Thus,

perturbations injected into the oscillator at this point must decay with time. Assuming that the frequency dependence of $Y_a(s)$ is negligible over the narrow band of interest, the oscillation will be stable provided the following statement holds:

$$\left(\frac{\partial |G_a(A)|}{\partial A} \Big|_{A=A_0} \right) \left(\frac{\partial B_r(\omega)}{\partial \omega} \Big|_{\omega=\omega_0} \right) + \left(\frac{\partial B_a(A)}{\partial A} \Big|_{A=A_0} \right) \left(\frac{\partial G_r(\omega)}{\partial \omega} \Big|_{\omega=\omega_0} \right) < 0,$$

where A and ω represent the amplitude and frequency of the oscillation, of which A_0 and ω_0 are the respective steady-state values. Since G_a is negative, it is important to note that the above expression involves the partial derivative of its magnitude. Implied in this equation is that slope of B_r must be positive, since it is the degradation of $|G_a|$ in VCNR-based resonant oscillator circuits that is responsible for limiting the amplitude of the signal. Hence the importance of monotonicity in $G_a(A)$: without this trait, oscillator stability cannot be guaranteed. The stability criterion means that when a voltage-controlled negative resistance active element has been used (as assumed in the discussion to this point), a parallel resonance must be chosen to reliably produce oscillations. Although this will be seen more clearly in the next section, an intuitive argument may be made for this case by noting that the primary variable about a parallel resonator is the voltage across it (as opposed to the current through each individual circuit element within it), thus a voltage-controlled negative resistance must be used with it.

If a current-controlled active element (ICNR) is used instead, one then needs to consider impedances, resistances, and reactances in place of their respective reciprocals. The valid set of criteria for steady-state oscillation with this type of circuit can be obtained by taking the dual of each of the previous relationships. In this case, the constraint that falls out of the stability consideration is that the current-controlled negative resistance element must be mated to a series resonance, for which the reactance has a positive slope with increasing frequency.

Words of caution are appended to this body of theory by Nguyen [17], who realized that some higher order effects could trip up the designer. Specifically, the presence of additional poles and zeros in the transfer function (often from parasitic elements) can make the basic steady-state oscillation criteria appear to be met from impedance or admittance plots, when in fact the circuit is stable. Another potential pitfall occurs when the

oscillation equations are satisfied at more than one frequency. This results in simultaneous oscillation at each of these frequencies, generating a signal with multiple spectral components. These possibilities can be addressed by simply adhering to good, basic circuit design practice. The designer should always investigate the operation of the circuit over a broad sweep of frequency and bias points to ensure desired operation. A plot of the impedance at the resonator port over frequency will indicate multi-oscillations if more than one zero-crossing of the phase lies within the negative resistance band (the frequency range over which the active element produces a negative resistance). Nguyen further recommends use of Nyquist or root locus plots (with bias current as the “gain”); in the latter case, the designer should verify that one set (and only one set) of complex conjugate poles are in the right-half of the s -plane, and that they remain comfortably there over a range of bias conditions. Transient (time-domain) simulations are also a valuable tool, one which usually provides clear indications of the existence of either circuit stability or multi-oscillation.

Once the circuit is known to oscillate, attention turns toward estimating the amplitude of the signal which is produced. Again, time-domain simulations provide a good indication, the accuracy of which is limited only by that of the large-signal device models. An iterative frequency-domain approach may also be used. The oscillation amplitude stabilizes where the energy supplied by the active element balances that dissipated in the resonator; for a voltage-controlled negative resistance, this occurs when the magnitude of the negative conductance equals the conductance of the parallel resonator to which it is joined. When a current-controlled device is used, oscillation builds until the magnitude of the negative resistance has fallen to the point where it is equaled by the parasitic resistance of the series resonator.

2.3 Resonator Circuit Models

Regardless of how the resonator used in conjunction with the negative resistance element is physically implemented, lumped inductor-capacitor (LC) circuit models are widely used to model their behavior. With the proper choice of component values, LC circuits are a valid method of modeling a wide variety of second-order systems. Loss is modeled by the addition of a resistor, and is often quantified in terms of the quality factor

(Q) of the second-order system. In such systems of the bandpass variety, Q is also equal to the center frequency of the filter's response divided by the half-power bandwidth of it. For a series RLC circuit, the quality factor expressed in terms of the circuit parameters is $Q = Z_0/R_s$, where Z_0 is the characteristic impedance of the LC pair ($Z_0 = \sqrt{L/C}$). With a parallel configuration, the reciprocal applies: $Q = R_p/Z_0$. For convenience in working with resonator circuits, the frequency response of a representative example ($f_0=1\text{GHz}$, $Q=5$) has been plotted in Figure 2-6 on the following page. The magnitude and phase, along with the real and imaginary components, of the impedance of a parallel RLC circuit appear in the left-hand column, with the associated graphs for admittance accompanying on the right. The series resonator is simply the dual of this case, for which the impedance plots are those in the right column, along with admittance on the left. Here it can clearly be observed (for oscillator stability purposes) that it is the susceptance of parallel resonator and the reactance of a series resonator which have positive slopes with frequency.

Often, the loss associated with an inductor and a capacitor is known, and it is useful to determine the quality factor of the resonator constructed from them. If the loss in each of the reactive components is modeled by a series resistance, the equivalent parallel resistance may be shown to be $R_p = Z_0^2 / (R_L + R_C)$, as indicated in the diagram below. The equivalency of the circuits holds for both frequency response (f_0 , Q) and for thermal noise generated by the resistive components. Alternatively, and more generally, the quality factor of the resonator may be expressed in terms of the Q of each component:

$$\frac{1}{Q} = \frac{1}{Q_L} + \frac{1}{Q_C}$$




Figure 2-5. Parallel resonator equivalent circuits using simple modeling of loss.

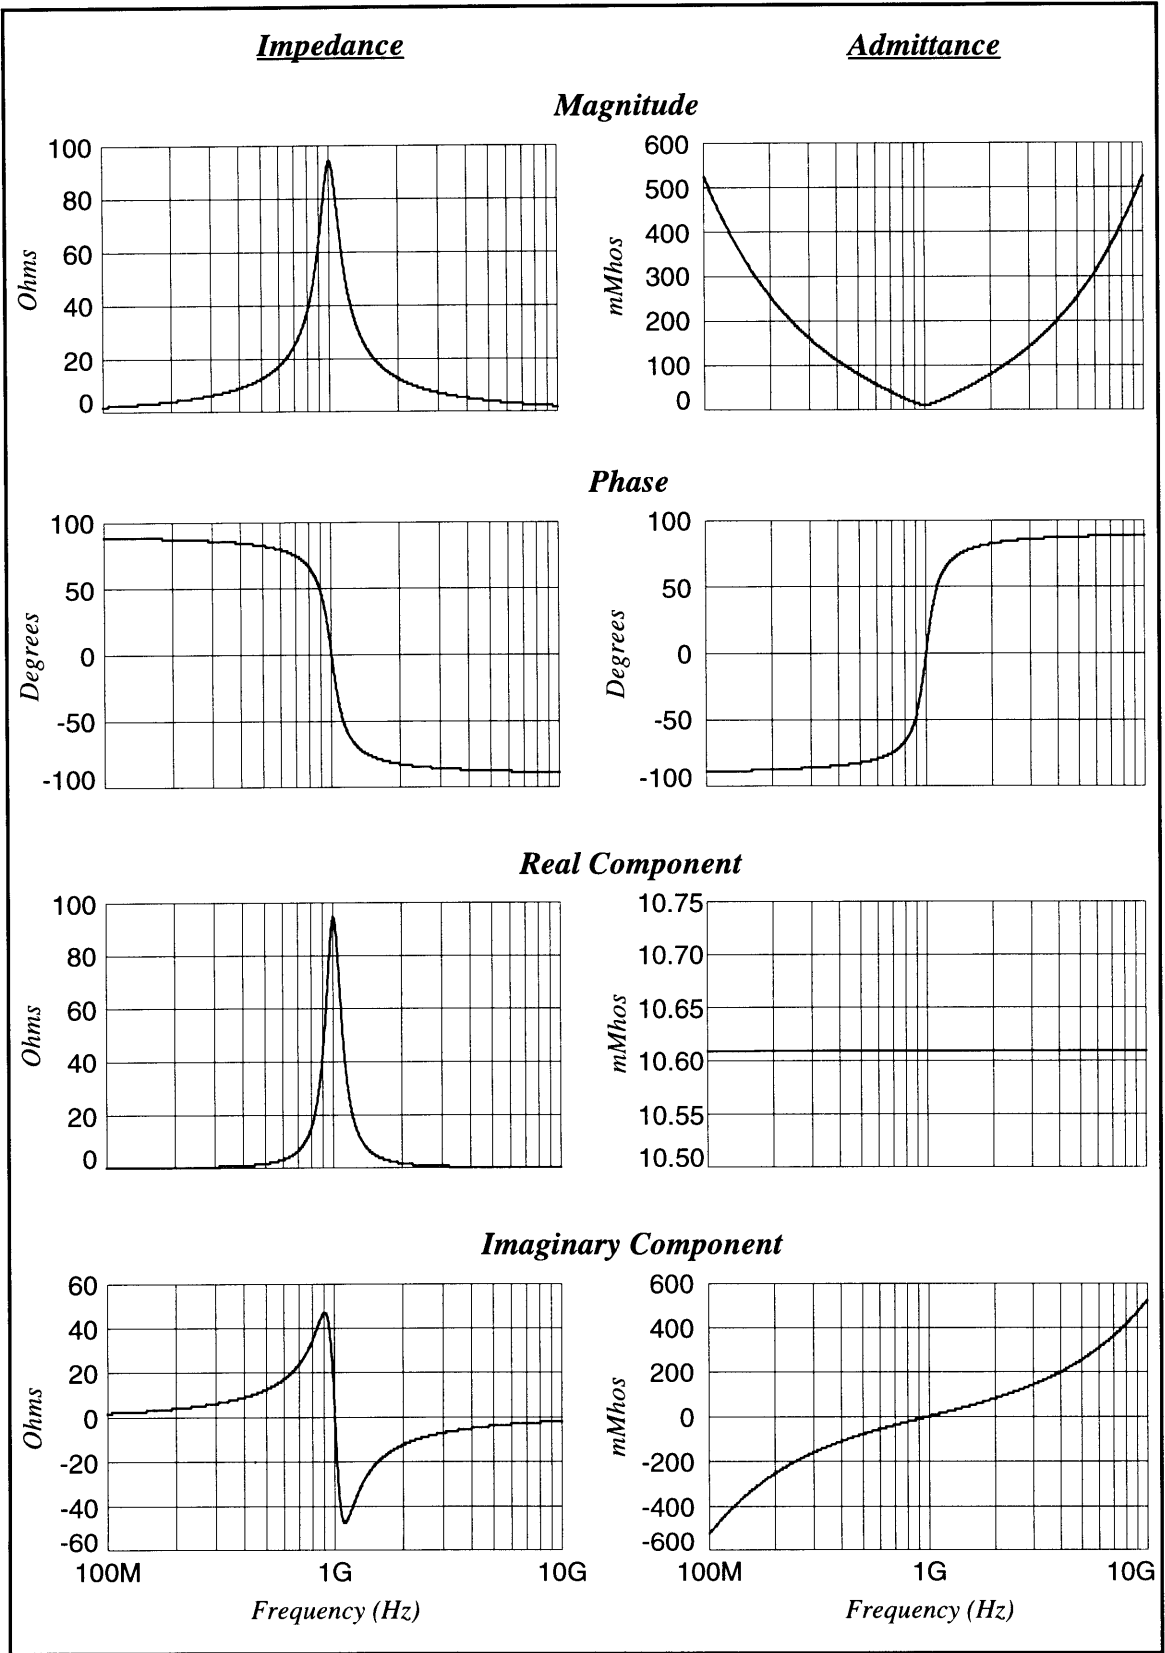


Figure 2-6. Impedance and admittance plots for a parallel resonator ($f_0=1$ GHz, $Q=5$).

2.4 Oscillator Noise Theory

As suggested by the discussion of resonator circuit models, even a strict adherence to the theory presented thus far will not yield an oscillator which provides an ideal impulse in the frequency spectrum. This spectrum of finite width results from noise in the oscillator circuit, and is shaped in conjunction with the resonator and circuit non-linearities. Oscillator noise is generally considered as being comprised of two components. The first of these is the spectrum resulting from time-varying changes in the amplitude of the oscillation, and is termed amplitude, or AM, noise. The latter piece of the spectrum is owed to random perturbations in the phase of the oscillatory signal, and has become labeled phase noise (or FM noise).

For wireless transmitter applications (particularly in a crowded frequency space), it is imperative to investigate the origins of the noise spectrum and its relationships to the circuit variables. Leeson [18] performed some of the original work in oscillator noise theory with a plausibility argument to explain the characteristics of the spectrum resulting from phase noise in a feedback oscillator. A diagram of the single-sideband oscillator spectrum is shown below in Figure 2-7, where the power spectral density relative to carrier power is plotted versus offset from the carrier frequency.




Figure 2-7. Leeson's oscillator phase noise model.

Leeson argued that for frequencies outside of the resonator corner frequency ($f_0 / (2Q)$), the resonator is effectively out of the circuit, leaving only the thermal (Johnson) noise power of the active gain stage, folded around the carrier (contributing a factor of

two), and made worse by the noise factor (F) of the active gain element. This sets the noise floor of the spectrum, and is shown relative to the power in the carrier. For frequency offsets inside of the $f_0/(2Q)$ corner, random fluctuations in phase at the input of the gain stage—resulting from voltage and current noise—are translated into frequency deviations by the phase-to-frequency transfer characteristic of the resonator (e.g., refer to Figure 2-6). Leeson reasoned that since the resonator is a second-order filter, the noise spectrum should correspondingly have a second-order roll-off. Finally, nearby the carrier, the oscillator will exhibit upconversion of flicker noise from the active elements used to implement the circuit, with the point of transition to the -30dB/decade slope being the device 1/f noise corner.

More than a decade later, Sauvage [19] demonstrated this view to be true quite generally, as did Lindenmeier [20] several years later, who then also verified that the effects on the oscillator spectrum due to amplitude noise were small compared with that resulting from phase noise. Using a generalized polynomial function to describe the non-linear gain element, Lindenmeier derived the power spectral density due to noise for the band characterized by a second-order roll-off (flicker noise and the thermal noise floor were not explicitly considered). Separating the components arising from amplitude and phase variations, and expressing them in terms of circuit design parameters, Lindenmeier achieved:

$$S_{AM}(\Delta f) = 10\log \left[\left(\frac{kTF}{P_o} \right) \frac{4\epsilon^2}{1 + 4\epsilon^2 Q^2 \left((\Delta f)/f_0 \right)^2} \right],$$

$$S_{FM}(\Delta f) = 10\log \left[\left(\frac{kTF}{P_o} \right) \left(\frac{1}{Q^2} \right) \left(\frac{1}{((\Delta f)/f_0)} \right)^2 \right],$$

which are in dBc/Hz (power in a 1 Hz bandwidth in decibels relative to the carrier power). Epsilon (ϵ) is a ratio of the relative change in oscillation amplitude over a given change in loading, and is generally on the order of unity. For large offsets from the carrier frequency (i.e. large Δf), the contributions from both noise components are equal. However, inside of the resonator's corner frequency, the spectrum due to amplitude variation levels off while that resulting from phase noise continues its second-order slope. Since this same

corner frequency also marks the entrance of the noise floor into the picture, the oscillator's phase noise dominates for all frequency offsets of interest.

Through this work, it is seen that the quality factor of the resonator is a big determining factor, since the phase noise spectrum is reduced by Q^2 . Power, however, also plays a significant role, entering the picture both through the noise figure of the oscillator gain stage, which is dependent upon the bias of the active devices, and also through the output power in the carrier (P_o) and its association to the DC power consumed in the VCO circuit. This relationship sets up a design trade-off between phase noise and power consumption, about which size of the active device(s) is the fulcrum. Placement of the "fulcrum" is a fundamental circuit design issue, and is discussed further in Chapter IV.

2.4.1 Oscillator Noise Matching

Optimization of the transistor geometry and bias to minimize noise figure is exactly the procedure which is followed for designing low noise amplifiers (LNAs) used in RF/microwave applications. With LNAs, the designer typically exercises one additional degree of freedom: the source impedance. A matching network is frequently used to transform the impedance presented to the LNA (e.g., 50Ω , real) to that which minimizes the noise figure of the active device(s) (i.e. Γ_{OPT}). Lindenmeier [20] realized that something akin to this can be carried out with oscillators.

An oscillator requires poles in the right-half of the s-plane (a negative resistance). These poles are usually created by placing positive feedback around an amplifier stage, such as demonstrated earlier in Figure 2-1. In these circuits, the feedback network carries some fraction of the output signal back around to the input, an action which is often modeled in circuit theory by a transformer. Within this representation, the source impedance presented to the amplifier is the output impedance of the oscillator gain stage divided by the square of the transformer turns ratio. By modifying the feedback (i.e. turns) ratio, a measure of impedance matching can be achieved. The ability to minimize phase noise in this fashion is consistent with the concepts of noise matching, and is a valuable tool for improving oscillator performance.

2.4.2 Other Phase Noise Analysis Methods

In his generally well accepted work, Lindenmeier modeled a transistor as a non-linear transconductance element, teamed it with a resonator, and derived a phase noise model dependent upon quality factor and noise figure—two linear circuit parameters. Recent thoughts that this first-order treatment may be optimistic, coupled with a desire to incorporate additional circuit elements and effects into the calculation, have led to a renewed push for methods of phase noise analysis and simulation. Once an accurate computational tool has been identified, numerical oscillator circuit optimization may be performed.

Some success has been reported using Kärtner's approach to calculating phase noise [21][22], but this method is not easy to implement, and is not readily available to the circuit designer. New commercial tools such as Compact Software's Microwave Harmonica have been imbued with phase noise analysis capability [23], but as of yet have not found widespread usage in published work. Another option is to use a time-domain simulator to perform a transient noise analysis, which is a traditional SPICE-like transient analysis in which current noise sources are appended in parallel with each noisy circuit element. Each "noise" source consists of a number of sinusoidal current sources with random amplitudes, phases, and frequencies. One commercially available product, Eldo, offers a feature to automate this analysis. Following this with a DFT makes for an interesting possibility, although the CPU load is necessarily heavy to obtain the required numerical accuracy from the transient simulator engine.

The discussion in this section has focused upon harmonic oscillators. Theory of noise in relaxation-based oscillators uses a different analytical tact, based on time domain jitter rather than frequency domain noise and shaping. For work in this area, the reader is referred to the oft-mentioned Abidi paper [2], or to two more recent efforts [24][25].

References

1. T. Aytur and B. Razavi, "A 2GHz, 6mW BiCMOS Frequency Synthesizer," *ISSCC Digest of Technical Papers*, International Solid-State Circuits Conference, Feb. 1995, pp. 264-265.
2. A. A. Abidi and R. G. Meyer, "Noise in Relaxation Oscillators," *IEEE Journal of Solid-State Circuits*, December 1983, pp. 794-802.

3. M. Thamsirianunt and T. A. Kwasniewski, "CMOS VCOs for PLL Frequency Synthesis in GHz Digital Mobile Radio Communications," *IEEE Custom Integrated Circuits Conference Proceedings*, 1995, pp. 331-334.
4. M. Soyuer and J. D. Warnock, "Multigigahertz Voltage-Controlled Oscillators in Advanced Silicon Bipolar Technology," *IEEE Journal of Solid-State Circuits*, April 1992, pp. 668-670.
5. J. Wu, "A Bipolar 1GHz Multi-Decade Monolithic Variable-Frequency Oscillator," *ISSCC Digest of Technical Papers*, International Solid-State Circuits Conference, Feb. 1990, pp. 106-107.
6. P. Alinikula, et. al., "Integrating Active Resonators for Wireless Applications," *Microwave Journal*, January 1995, p. 106.
7. K. W. Kobayashi and A. K. Oki, "A Novel Heterojunction Bipolar Transistor VCO Using an Active Tunable Inductance," *IEEE Microwave and Guided Wave Letters*, July 1994, pp. 235-237.
8. A. Ali and J. L. Tham, "A 900MHz Frequency Synthesizer with Integrated LC Voltage-Controlled Oscillator," *ISSCC Digest of Technical Papers*, International Solid-State Circuits Conference, Feb. 1996, pp. 390-391.
9. M. Soyuer, et. al., "A 3V 4GHz nMOS Voltage-Controlled Oscillator with Integrated Resonator," *ISSCC Digest of Technical Papers*, International Solid-State Circuits Conference, Feb. 1996, pp. 394-396.
10. J. Craninckx and M. Steyaert, "A CMOS 1.8 GHz Low-Phase-Noise Voltage-Controlled Oscillator with Prescaler," *ISSCC Digest of Technical Papers*, International Solid-State Circuits Conference, Feb. 1995, pp. 266-267.
11. K. K. Clarke and D. T. Hess, *Communication Circuits: Analysis and Design*, Addison-Wesley, Reading, MA, 1971, Chapter 6.
12. G. Gonzalez, *Microwave Transistor Amplifiers: Analysis and Design*, Prentice Hall, Englewood Cliffs, NJ, 1984, Chapter 5.
13. S. M. Sze, *Physics of Semiconductor Devices*, second edition, John Wiley & Sons, New York, 1981, Chapters 9-10.
14. J. J. Carr, *Secrets of RF Circuit Design*, TAB Books, New York, 1991, Chapter 20.
15. K. Kurokawa, "Some Basic Characteristics of Broadband Negative Resistance Oscillator Circuits," *Bell System Technical Journal*, July-August 1969, pp. 1937-1955.
16. E. C. Niehenke, "RF and Microwave Oscillator Design," from a short course entitled *Applied RF Techniques I*, Besser Associates, Los Altos, CA, 1996.

17. N. M. Nguyen and R. G. Meyer, "Start-up and Frequency Stability in High-Frequency Oscillators," *IEEE Journal of Solid-State Circuits*, May 1992, pp. 810-820.
18. D. B. Leeson, "A Simple Model of Feedback Oscillator Noise Spectrum," *Proceedings of the IEEE*, February 1966, pp. 329-330.
19. G. Sauvage, "Phase Noise in Oscillators: A Mathematical Analysis of Leeson's Model," *IEEE Transactions on Instruments and Measurement*, December 1977, pp. 408-410.
20. H. K. Lindenmeier, "Noise Matching Techniques in Transistor Oscillators," 1980 *IEEE Circuits and Systems Symposium*, pp. 1052-1055.
21. V. Güngerich, et. al., "Phase Noise Reduction of Microwave Oscillators by Optimization of the Dynamic Behaviour," *1994 IEEE MTT-S Digest*, p. 953-956.
22. W. Anzill, et. al., "Phase Noise Minimization of Microwave Oscillators by Optimal Design," *1995 IEEE MTT-S Digest*, p. 1565-1568.
23. U. L. Rohde and C.-R. Chang, "Analysis and Optimization of Oscillators for Low Phase Noise and Low Power Consumption," *RF Design*, March 1995, p. 70.
24. T. C. Weigandt, et. al., "Analysis of Timing Jitter in CMOS Ring Oscillators," 1994 *IEEE Circuits and Systems Symposium*, pp. 27-30.
25. B. Kim, et. al., "PLL/DLL System Noise Analysis for Low Jitter Clock Synthesizer Design," 1994 *IEEE Circuits and Systems Symposium*, pp. 31-34.

III. Monolithic Inductors

Inductors, the “forgotten” elements in much of conventional integrated circuit design, find wide application in circuits designed to work at the upper RF and microwave frequency ranges. Impedance matching to the external world to maximize power transfer and to minimize distributed effects cannot, in general, be accomplished with resistances and capacitances alone. Inductances are also valuable in constructing frequency-selective filters, isolating DC supplies from RF signal energy, tuning amplifiers for narrowband or low noise applications, and promoting stability in gain stages. For signals approaching (and in) the microwave portion of the spectrum, usable values of inductors fall below 10nH—values which may be achieved monolithically. Inductance is only one aspect, however; although recent research efforts have yielded some improvements [1][2], the capabilities afforded by current semiconductor technology are not adequate for all needs. Whether a monolithic solution is feasible depends on each usage.

3.1 System Considerations for the Inductors

Earlier chapters have expressed that the need for low power and good phase noise in an oscillator translates into a high quality (high Q) resonator. Waveguide (e.g., microstrip) resonant elements find popular application at higher microwave frequencies, but are too unwieldy at the desired 1.8GHz to be efficiently used on chip. Acoustic resonators are a possibility, but entail difficult IC manufacturing issues which need to be addressed; thus LC tanks remain a more viable option.

Semiconductor processing techniques lend themselves well to manufacturing good monolithic capacitors. Quality factors above 30 are not uncommon, and values as high as 80 have been reported with standard interconnect technology [1]. Most of the loss in integrated, lumped-element resonators thus accrues due to the on-chip inductors which have only been available (in typical silicon processes) with Q's approaching 5 [3][4]. Several efforts have tried to remedy this situation by using active, inductance-simulating circuits [5]-[8], but generally have resulted in high power consumption. For low power applica-

tions, passive inductors integrable onto silicon substrates are preferable. Along with these criteria, structures with inductances in the 1 to 10nH range, increased quality factors, and self-resonant frequencies of at least 4 to 5GHz are needed for L-band circuits.

3.2 Monolithic Inductor Test Wafer

To help achieve the desired inductor and circuit performance, a set of monolithic structures were designed and fabricated in MTL. P-type silicon wafers, taken from the 14 to 24 Ω -cm bin, were chosen for the starting material as being representative of a typical, high-speed, bipolar or BiCMOS process. A top-level metal was simulated with a deposition of 1 μ m thick AlSi metallization onto 2.15 μ m of field oxide, and then patterning the metal with a plasma etch. For timeliness of design and processing, only one mask was used, a plot of which is shown in Figure 3-1 on the following page.

The spiral inductors on this test wafer were drawn with 3, 4, and 5 turns, and in both square and octagonal shapes. The width and spacing of the metal lines composing these structures were varied in terms of their ratio in such a way that the total spiral area was kept nearly constant for a given number of turns (i.e. the metal pitch remained the same). Three width-to-spacing ratios were constructed; the base line consisted of 15 μ m wide conductors and 15 μ m of space between them, while “thick” utilized a 20 μ m/10 μ m ratio, and “thin” was the opposite—10 μ m/20 μ m. Since only one metal level was available, a bond wire was used to jumper the center point of the spiral out to a co-planar waveguide (microwave) probe pad in a one-port configuration. Although the presence of this jumper does change the measurements, the wire was characterized individually and its effects correspondingly backed out of the spiral inductor data.

In addition to the spirals, bond wire inductors were also built in one-port and two-port configurations. These rely upon the parasitic elements of wire, which are dominated by an inductive component for IC-sized bond wires at moderate microwave frequencies and down through RF. For manufacturing consistency, it is desirable to have the wires which form the inductors jump from one bond pad to another on the same integrated circuit die, as opposed to bonding from a die to the lead frame. Since the value of the inductance in the wire is primarily a function of the horizontal distance between bond pads, good tolerances can be achieved with the pad-to-pad bonds [9].

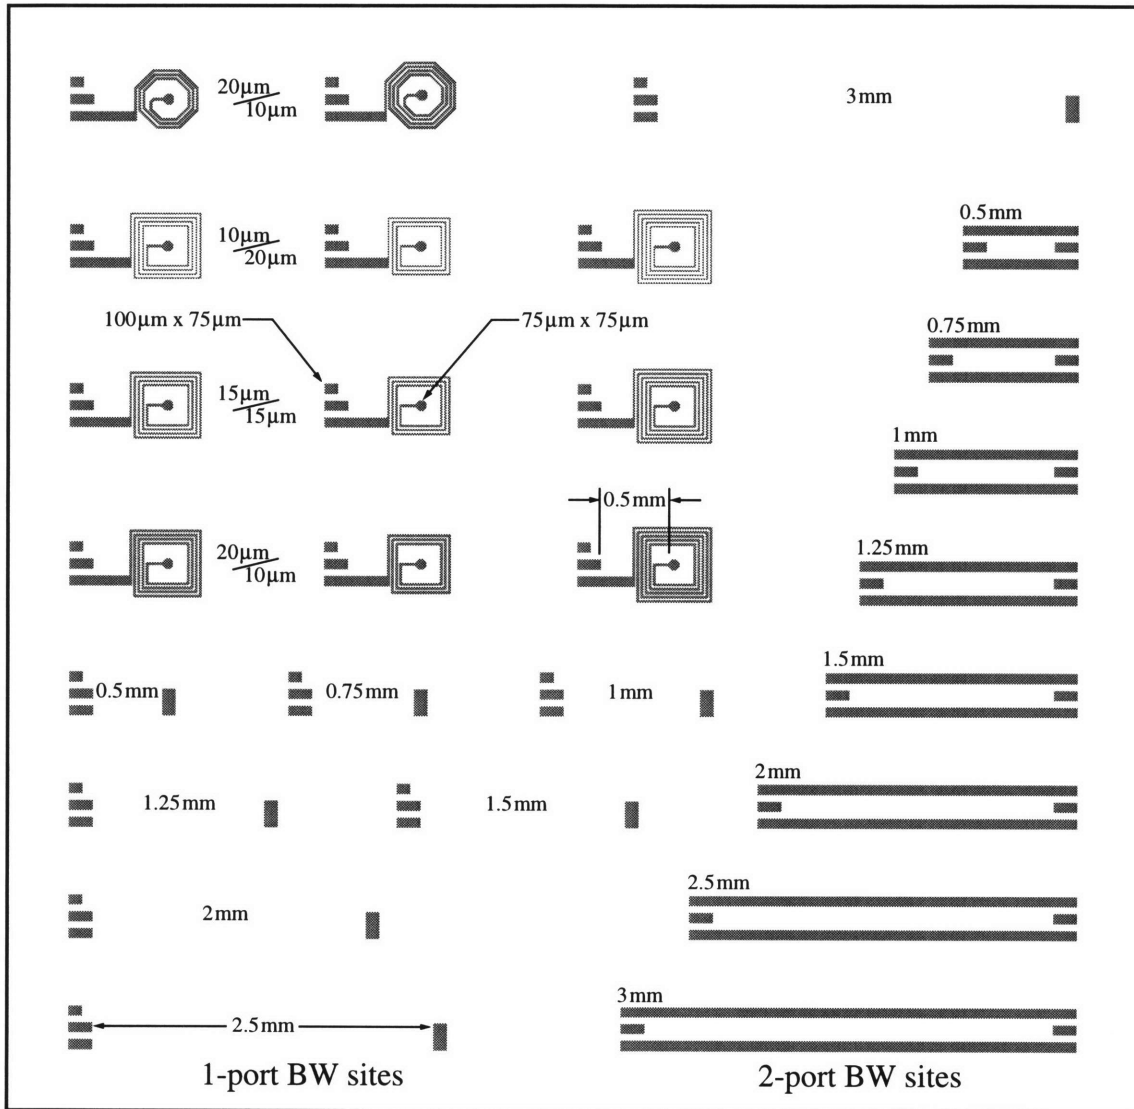


Figure 3-1. Metal mask pattern for monolithic inductor test wafer.

The two-port bond wire inductors consisted of a single wire bonded between pads set apart by a distance which was varied between 0.5mm and 3mm and encased by coplanar waveguide (CPW) probe pads. The one-port structures were built from two parallel bond wires as shown in Figure 3-2. One wire is bonded from a signal pad out to a “shorting bar” of metal, and the other runs back from this bar to the ground pad of the CPW probe. Both gold and aluminum bond wires of 1 mil ($25.4\mu\text{m}$) diameter were tried, along with both ball and wedge bonds. This latter type of bond offers a significantly smaller contact patch, but results in a portion of the wire at either end being relatively near the lossy substrate.




Figure 3-2. Photo of one-port bond wire inductors (1.25 mm and 2 mm).

3.3 Monolithic Inductor Wafer Test Methodology

As mentioned several times previously, each of the inductor structures were linked to co-planar waveguide probe pads to allow for accurate high frequency, on-wafer characterization. Cascade Microtech's K-band (up to 40GHz) WPH-series probes were used along with semi-flex cable to connect to a Hewlett Packard 8510C network analyzer, and the setup was calibrated using the standard open/short/ 50Ω load/thru-line methodology. Scattering matrix parameters were then measured and logged for the inductors, from which the desired information could be extracted. Although linear elements are being characterized, the source was set at a power level of 0dBm (1 mW)—representative of signal levels in the VCO—to prevent second-order effects from skewing measurements away from actual conditions of usage.

For the one-port inductors, s_{11} maps directly to an impedance (a conversion performed by the HP8510) which can be easily calculated or read from a Smith chart. The situation for the two-port bond wires is only slightly more involved, wherein s_{11} represents the impedance of the inductor in series with a 50Ω resistance (the load resistor). On a Smith chart, the plot of s_{11} versus frequency is the same as for the one-port, except that it is referenced to the 50Ω circle instead of the unit (0Ω) circle. The transmission parameter

data ($s_{21}=s_{12}$ in the case of passive elements) is also meaningful in this context as the frequency response of the inductor when used in a matched circuit. This information is useful for filter design, and also aids the distillation of physical effects from the experimental results.

The measured s-parameters can be directly input into some circuit simulators (e.g., Libra) to reproduce the operation of the inductor, or a lumped element model may be fit to the data to achieve the same functionality in SPICE engines. Important figures of merit can also be extracted from the measurements to compare and evaluate the structures, with the most important numbers in this application being inductance, quality factor, and self-resonant frequency. The value of the inductance provided by a structure can be calculated from the imaginary part of its impedance for each frequency point. Similarly, the quality factor (Q), defined as the ratio of energy stored in an element to energy dissipated by it, is equal to the magnitude of the imaginary part of the element's impedance divided by the real part. It is important to note that for these inductor structures, in general, $Im\{Z_L\} \neq \omega L$, as parasitic capacitances begin to roll off this value in the frequency range of interest. At the frequency where the reactance of these capacitors equals that of the inductance, the impedance is entirely real. This is the frequency of self-resonance, the point at which the structure stops behaving as an inductor.

3.4 Inductor Characterization Results

After fabrication was completed, four samples of each inductor were bonded with wire so that an expedient survey of the structures could be conducted. Following a comparison of the parameters and frequencies of interest, a few of the more promising inductors were chosen for more detailed characterization. The initial data also allowed some conclusions regarding bond wire material, wire diameter, and bond type, which were useful in framing the “fine resolution” experiment.

A summary of the survey findings is depicted on the following pages for both the spirals (Figure 3-3) and the bond wire inductors (Figure 3-4). The effects of the jumper wire used to access the inside of the spiral were estimated from two-port bond wire measurements and then backed out of the data presented in Figure 3-3. This graph has been divided into sections by the number of turns in the spiral: the inductance (L) and quality

factor (Q) of the 3-turn, 4-turn, and 5-turn structures are plotted from left to right. The key along the horizontal axis of the plot indicates that a given point corresponds either to an octagonal shape or to a square layout with the metal line width as shown. Both of the octagons were drawn with $20\mu\text{m}$ metal width and $10\mu\text{m}$ spacing as described earlier in Section 3.2 (i.e. the “thick” line width and spacing).




Figure 3-3. Survey of MTL spiral inductors at 1.8GHz.

It is interesting to note that for a square spiral, the inductance is fairly constant for a given number of turns and metal pitch, but the quality factor is improved as the metal width to spacing ratio is increased. This is consistent with the findings of S. Chaki, et. al. [2], and can result in a significant improvement for smaller-valued inductors. Another useful detail borne out in the above graphs is that the octagonal structure exhibited an improved Q when compared to its square brethren of the same number of turns. The inductance was also diminished (by about 15%), but the increase in Q (20% to 30%) overshadowed this drop. In general, the more circular the structure, the greater the quality factor will be for a given inductance, as resistive corners are alleviated and finally eliminated as a circle is approached. This fact—that a significant amount of the series resistance in a

spiral occurs in the corners—is also the motivation behind the design guideline to leave open the middle area of the spirals: tight turns introduce many corners without adding much inductance. This latter result stems from the nature of the inductance, which arises due to mutual coupling between parallel lines of some length. If the length is short, so too is the value of the inductor.

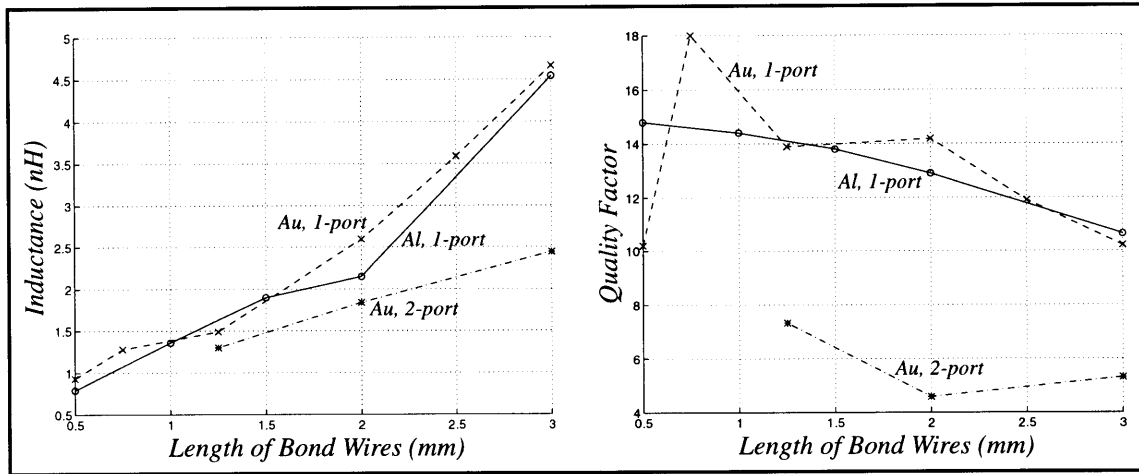


Figure 3-4. Survey of MTL bond wire inductors at 1.8GHz.

The inductance and the quality factor for the bond wire inductors are plotted above against the length of the bond wires used, with one-port structures comprised of two lengths of wire, and two-port inductors made of one. The data suggest a very nearly linear relationship between inductance and length, with no indication of material dependence at 1.8GHz. Since the value of the two-port inductor nearly equals that of the one-port constructed with the same lengths of wire, one may infer that mutual coupling exists between the parallel wires which has a deleterious effect upon inductance. Additional measurements and calculations could be carried out to isolate the self-inductance terms from the mutual coupling, but attaining this distinction was not a goal of the experiment.

Although difficulties in achieving consistent bonds with gold wire on the AlSi metallization resulted in some measurement scatter, it remains evident that gold wire did not provide a higher Q than did aluminum. This suggests that the primary loss mechanism in the bond wire inductors is capacitive coupling into the substrate, an assertion which is further supported by comparison of measured data to values computed by FastHenry. (FastHenry is an electromagnetic simulation tool developed at MIT which extracts induc-

tance and resistance parameters from arbitrary 3-D shapes, but does not account for capacitance.) Simulations using FastHenry predict quality factors of around 30 for the bond wire inductors at 1.8GHz, which is optimistic by more than a factor of two. The values of series inductance extracted by the simulations were also significantly higher than the measured reactances would indicate, again pointing to capacitive loading as rolling off this number.

All of the bond wire inductors performed very similarly with one exception: the two-port structures exhibited a much lower Q —numbers low enough, in fact, to be achievable by the spiral elements which were characterized. The lower values can be expected since, with the one-port configuration, one terminal of the bond wire network is placed at an AC ground, a move which shorts a parasitic-laden terminal of the device. Even so, the two-port inductors, along with the rest of the bond wire structures, achieved self-resonant frequencies (f_{SR}) above 20GHz (where the frequency sweep ended) until the longest wire lengths were reached. The longer wires made observable that f_{SR} for the gold one-ports was marginally lower than that for the aluminum, which brings up the last distinction between these inductors. This reduction in f_{SR} is not so much attributable to the element constituting the wires, but rather to the way in which the wires were attached to the metal. Both ends of each aluminum wire were attached via wedge bonds, while one bond for each of the gold wires was a ball bond (the equipment at MTL was incapable of terminating wires in ball bonds at both ends). Thus the results seem to indicate that the smaller contact area of the wedge overcomes any disadvantages which may accrue due to the non-perpendicular fashion in which this bond attaches the wire to the die.

The metal comprising the spiral inductors is both more resistive and closer to the lossy silicon substrate than that for the bond wires. These factors conspire to restrict the spiral devices to functioning at much lower frequencies, as is depicted in Figure 3-5 on the subsequent page. Coming somewhat as a surprise, f_{SR} for the spirals seems to be more closely linked to the total area occupied by the spiral, rather than the total area of metallization within it. For a square-shaped device of a given number of turns, f_{SR} remained constant, independent of the width to spacing ratio. As with Q , the octagonal structures exhibited a significant increase in f_{SR} over their square counterparts with the same number of turns.




Figure 3-5. Self-resonant frequency data for MTL spiral inductors.

Self-resonant frequency is a common metric for inductors, analogous in many ways to the unity current gain frequency (f_T) of an active device. A large quality factor is generally desirable with an inductor, but the value of Q goes to zero at f_{SR} . Similar to the way in which transistors need to be used in conventional IC design well away from f_T , an inductor should not be used at frequencies approaching f_{SR} . To be sure, such a use will result in a very low Q . Perhaps more importantly, however, is that near f_{SR} , higher-order parasitic effects which are not accurately modeled may predominate, potentially resulting in localized pockets of capacitive behavior (i.e. where $\frac{d}{df}(Im\{Z\}) < 0$) and an unpredictable phase response.

3.4.1 Detailed Data on Bond Wire Inductors

The survey of the monolithic inductor structures revealed several important findings: the bond wire inductors—particularly the one-port devices—outperformed the spirals, gold wires offered no discernible advantage, and the wedge-bonded aluminum wires provided the most repeatable results. In addition, the initial data set allowed for a cali-

bration of the bond wire lengths that are required to achieve the inductor values desired for the VCO circuitry. Although the data indicate that smaller-valued inductors yield higher quality factors, a balance needs to be struck for a resonator since capacitors also tend to perform poorly and consume significant die area as they become large. For the VCO application, inductances in the vicinity of 2nH were looked at with this trade-off in mind. Based upon the survey results, one-port inductors with aluminum bond wires were chosen in the lengths of 1.25mm, 1.5mm, 2mm, and 2.5mm for more detailed characterization. Ten additional structures in each of these four configurations were bonded and measured; the results at 1.8GHz are presented below in Figure 3-6 (inductance) and Figure 3-7 (Q).

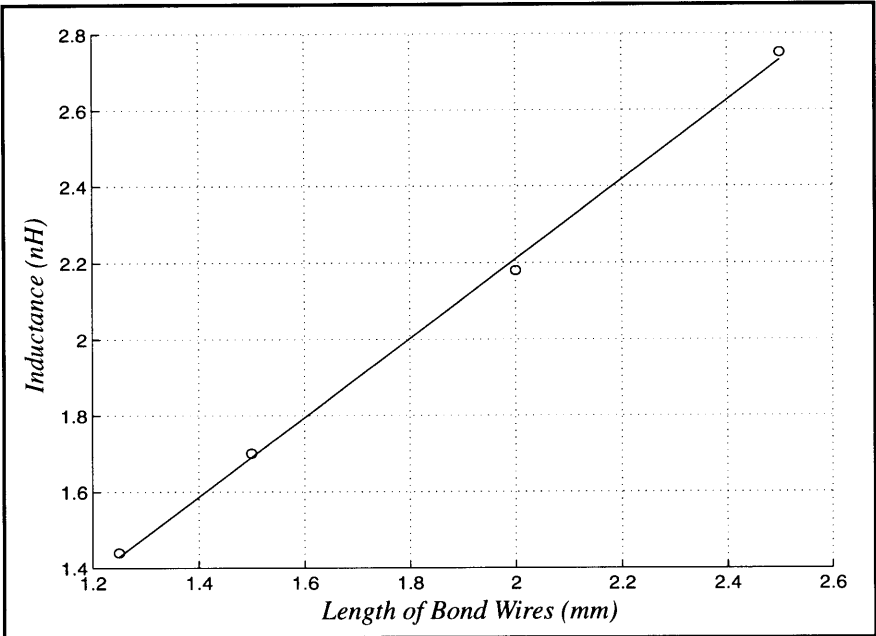


Figure 3-6. Inductance of MTL one-port bond wire inductors at 1.8GHz.

As before, the inductances of the one-port bond wire structures closely fit a linear model, lending further support to the premise that the inductance originates primarily from the horizontal lengths of wire. The data fit the line:

$$L = 0.13 + 1.04 (\text{length}),$$

where the length is given in mm and the inductance is in nH. These values are consistent with the work of J. Craninckx and M. Steyaert [9], in which numbers of approximately 1 nH per mm were achieved for similar devices.

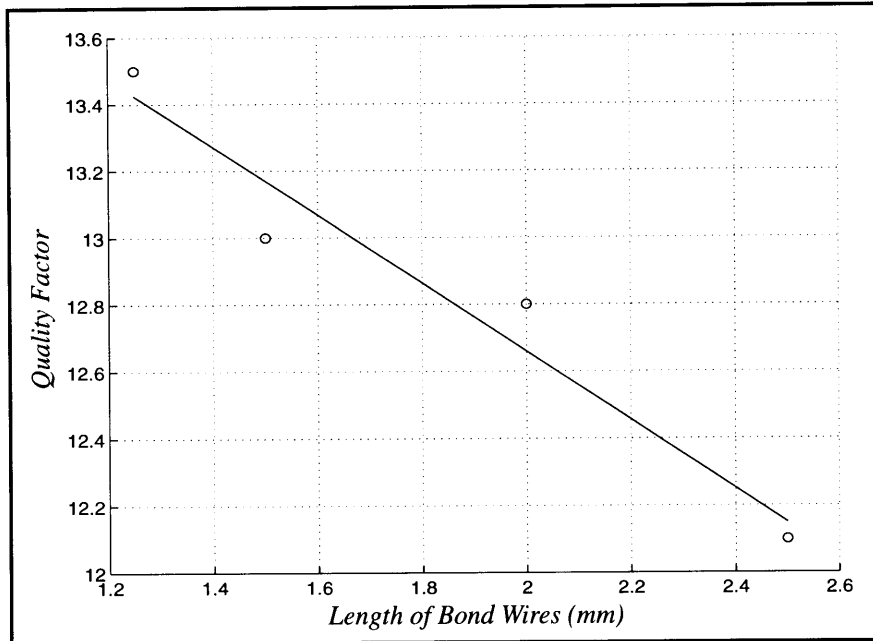


Figure 3-7. Quality factor of MTL one-port bond wire inductors at 1.8GHz.

Slightly increased scatter about a linear approximation was prevalent in the quality factor data, but the average value at each length of wire was still within one standard deviation of the norm. The interpolation model here is:

$$Q = 14.7 - 1.02 (\text{length}),$$

with length in mm. Although one needs to exercise caution when attributing physical significance to a statistical fit, the low value of the y-intercept—the “theoretical” Q with zero-length bond wires—seems to support the notion that a large portion of the loss in this structure is occurring at the bonding sites.

Inductance and quality factor have been addressed, but one issue that remains is that of repeatability. Some doubt has been raised regarding bond wire inductors [10], but this has largely concerned bonds from the silicon die to the lead frame. In this study of inductance elements constructed with a manual bonding machine, the 1.5mm, 2mm, and 2.5mm wires yielded 3σ tolerances on the inductance of $\pm 12\%$ (the variation in the shorter 1.25mm wires was slightly greater). This is comparable to the absolute accuracy achievable for monolithic capacitors, and these numbers could improve with automated equipment adapted for this operation. Since bond wire inductors are added following the

conclusion of fabrication, there may exist an additional possibility: gauging the capacitance in the resonator at wafer probe. This could allow a measure of course trimmability in the oscillator center frequency by alteration of the nominal bond length (i.e. the layout of a chip could allow for multiple bond sites).

3.4.2 Lumped Element Model of S-parameter Data

Although the linear regression analyses of the previous section could be utilized to interpolate results for any length of wire (and could even form the basis of a scalable model), the most accurate “model” is the measured data itself. The 2mm structure was selected for use in the VCO circuit due to its proximity to the desired inductance value. The s-parameter data for each sample of this length were inspected for outliers and then averaged at each frequency point. In order to reproduce these measured s-parameters in SPICE simulations, a circuit model was developed which contained an inductor, some series resistance, capacitive coupling elements both along the wire and to the substrate, plus some resistive losses within the substrate. This circuit was connected as a one-port network, and then the HSPICE optimizer was set upon it.

A least square error optimization drove the fringing capacitance (between the parallel bond wires in this case) term to zero. This is not entirely surprising since the wires are separated by a fair distance (125 μ m). A “T” model for the bond wire was then tried, but its center leg was similarly optimized away by the HSPICE algorithms. Left behind was the frequently used circuit model depicted on the following page in Figure 3-8, where all of the parasitic capacitance is lumped in with the bond pads at either end of the wire. This model fits the data well, as shown by the comparison of the modeled to measured s-parameters in Figure 3-9. Although some small, higher frequency excursions—which were consistently observed for the inductors—are not handled well by the model, its validity remains strong within the range of interest.

Even though there were no intentions of creating a broadly scalable device representation, nor desires of being able to extrapolate from it, the resulting circuit does seem to embody a good deal of physical intuition. The value of the series resistance fit to the data is comparable to that extracted by FastHenry calculations at 1.8GHz (where the skin effect is significant for bond wires). The capacitors to substrate are about what one would expect for top-level metal only bond pads, and the substrate resistance is on par with that

witnessed for capacitive structures in more detailed EM simulations as well as in device characterization data pertaining to substrate ties.

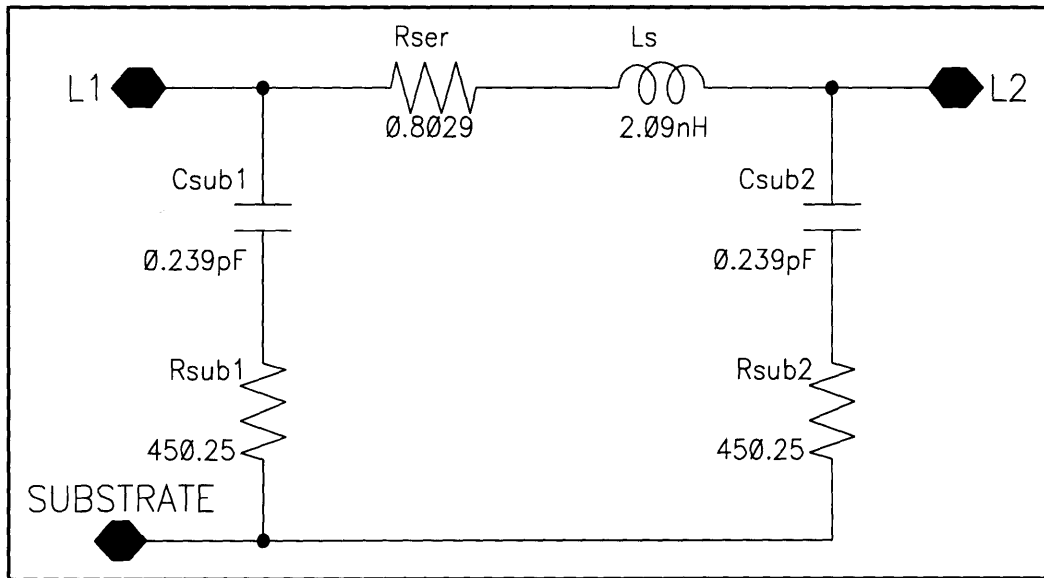


Figure 3-8. Lumped element model of the 2mm one-port bond wire inductor.

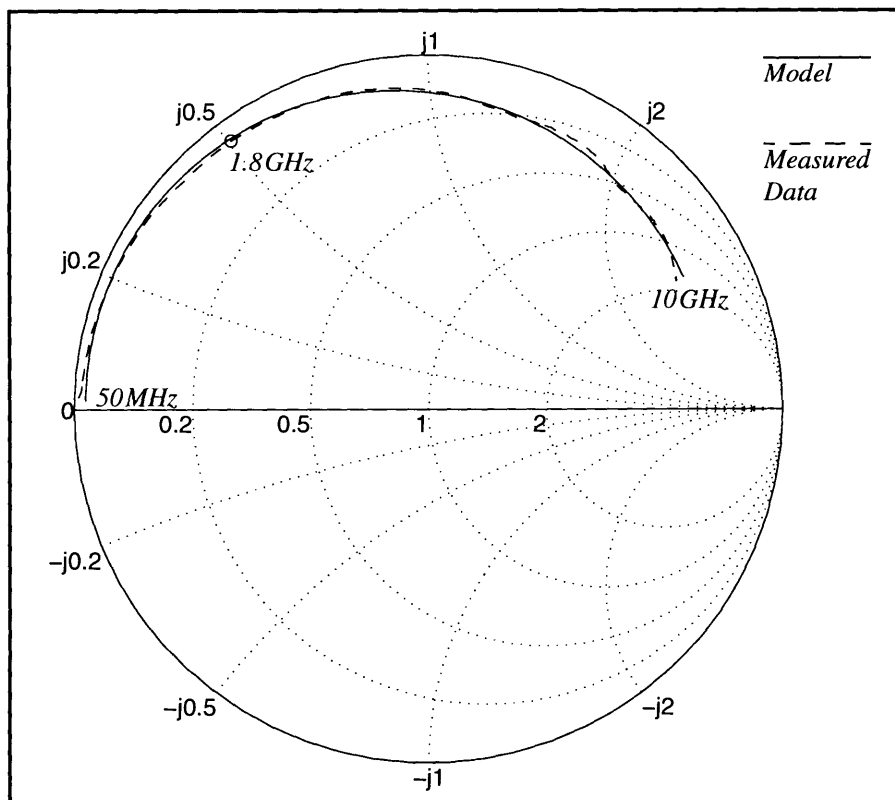


Figure 3-9. Comparison of SPICE model to s-parameter data.

3.5 Other Possibilities

A number of additional thoughts may come to mind regarding the improvement of monolithic inductor structures such as those employing bond wires. A high resistivity substrate reduces the effects of coupling capacitances, thereby increasing the quality factor of inductors built upon it. However, this tact is incompatible with production silicon wafer processes because it requires expensive float-zone crystal wafers. Thicker dielectrics place the inductive elements further above the substrate, reducing the value of the metal to substrate capacitance and thereby increasing Q . Again, this benefit is provided to any metal structure built upon the deeper insulator, but the designer is limited in this regard by the available process.

Thicker metal helps the spiral inductors [1], but moving to larger diameters will not likely improve the bond wire elements because the skin depth (about $1\ \mu\text{m}$ in aluminum at 2GHz) is a small fraction of even the narrowest wires. Although, in this study, gold did not appear to provide any advantages, this may be a consequence of the fact that the bonds were not as good as those achieved with aluminum. In any event, for the bond wire inductors, the potential for improvement due to a more conductive medium is tempered by the issue of skin depth, which is inversely proportional to the square root of conductivity ($\delta \propto 1/\sqrt{\sigma}$). Hence, in moving to gold wire from aluminum, the reduction in the series resistance term will be at most 10%.

Another strategy to improve resonator quality factors is to contain the energy in non-lossy dielectrics. Using a lower metal level to shield an inductor structure from the silicon substrate can improve Q , but unless the dielectric is very thick, the inductance is reduced to values too low to be worthwhile. So while this rules out the use of ground planes for spiral inductors built with standard interconnect technology, the possibility remains for the bond wire devices. (Experimental data indicate that the inductance accrues mostly to the horizontal length of wire between bonds, while much of the loss occurs at the pads.) The test wafer designed for this experiment did not have the required second level of metallization, but an inductor structure with a ground plane beneath it has been implemented on LMVCO chip silicon. The effect of the ground plane may then be evaluated based on the performance of this inductor compared to a control device without the lower level of metal.

References

1. J. N. Burghartz, M. Soyuer, and K. A. Jenkins, "Microwave Inductors and Capacitors in Standard Multilevel Interconnect Silicon Technology," *IEEE Transactions on Microwave Theory and Techniques*, Jan. 1996, pp. 100-104.
2. S. Chaki, et. al., "Experimental Study on Spiral Inductors," *1995 IEEE MTT-S Digest*, pp. 753-756.
3. R. Duncan, et. al., "A 1 GHz Sinusoidal Oscillator," *IEEE Custom Integrated Circuits Conference Proceedings*, 1995, pp. 91-94.
4. N. M. Nguyen and R. G. Meyer, "A 1.8 GHz Monolithic LC Voltage-Controlled Oscillator," *IEEE Journal of Solid-State Circuits*, March 1992, pp. 444-450.
5. P. Alinikula, et. al., "Integrating Active Resonators for Wireless Applications," *Microwave Journal*, January 1995, p. 106.
6. K. W. Kobayashi and A. K. Oki, "A Novel Heterojunction Bipolar Transistor VCO Using an Active Tunable Inductance," *IEEE Microwave and Guided Wave Letters*, July 1994, pp. 235-237.
7. B. Hopf, et. al., "Coplanar MMIC Active Bandpass Filters," *1994 IEEE MTT-S Digest*, pp. 1183-1185.
8. P. R. Nuytkens, et. al., "A Monolithic Tunable High-Q InAlAs/n+-InP-HFET Resonator," *Fifth International Conference on InP and Related Materials Proceedings*, April 1993, pp. 389-392.
9. J. Craninckx and M. Steyaert, "A CMOS 1.8 GHz Low-Phase-Noise Voltage-Controlled Oscillator with Prescaler," *ISSCC Digest of Technical Papers*, International Solid-State Circuits Conference, Feb. 1995, pp. 266-267.
10. P. R. Gray and R. G. Meyer, "Future Directions in Silicon ICs for RF Personal Communications," *IEEE Custom Integrated Circuits Conference Proceedings*, 1995, pp. 83-90.

IV. 1.8GHz (L-band) Monolithic VCO Design

It has often been said that the best way to get a functional oscillator circuit is by trying to design an amplifier. While it may be true that embedding high gain stages in a medium replete with parasitic reactances opens the door to uninvited oscillations, increasingly accurate modeling and simulation, along with appropriate design practices, have generally allowed the intended purpose to carry the day. Similarly, with VCOs, careful attention to design guidelines and simulation of parasitic elements and their effects on phase and feedback goes a long way toward minimizing uncertainty in the final product. The theoretical underpinnings of oscillator design were discussed in Chapter 2, including the important concepts of start-up, stability, and noise performance. Being able to effectively deal with these issues releases the designer from the constraints of existing circuits, allowing for creative exploration of new solutions to meet new challenges, while mitigating the risks of trying something different.

4.1 System Considerations for the VCO

Previous chapters have provided the background of an ultra-low power wireless sensor with a high performance 1.8GHz transmitter. In the proposed architecture, the data to be transmitted modulates the frequency synthesizer, within which the VCO produces the RF spectrum to be output to the antenna. Since the oscillator's phase noise becomes a part of the transmitted spectrum, a low noise design is critical to system operation. However, for the system to be useful, it must survive on a limited supply of battery power; hence minimization of its consumption is also an important goal. To meet these criteria, and in keeping with overall system size targets, it was argued in Chapter 2 that a fully monolithic, silicon-based, harmonic oscillator is the desired approach.

One of the primary challenges to designing high-frequency circuits in a silicon substrate is its (relatively) high conductivity (for a dielectric). The first problem this poses is that silicon makes a poor dielectric in which to store or propagate electromagnetic fields. The loss that occurs in the substrate makes achieving high Q inductors and resonator structures a formidable challenge—an issue covered in Chapter 3. Another difficulty

with silicon is that it offers limited isolation between circuits, particularly a problem in highly integrated systems. Thinning the wafers helps, as does liberal use of substrate ties in the layout (to shunt substrate noise to a low impedance ground or supply), including grounding of the back side of the die when possible. Adding deep trench isolation around devices is another beneficial step, but is not available in all processes. From a circuit design standpoint, utilization of fully differential topologies to reduce the effects of common-mode noise—which couples into signal paths from the substrate—should be strongly considered wherever feasible.

In addition to striving for the noise and power targets set forth in the introductory chapter, the VCO transfer function and interfaces must be designed in conjunction with the transmitter. The critical parameters of an oscillator transfer's characteristic are gain, linearity, and tuning range. VCO gain, typically labeled κ_v , relates the change in the frequency of the output signal for a given change in the input control voltage. Linearity, in this context, refers to the degree to which the gain remains constant over the tuning range. This range of output frequencies that can be produced by the VCO is usually limited by the extents of the signals which may be applied to the tuning element (along with the gain achievable by this element).

As used in the synthesizer of the wireless sensor, the input signal to the VCO will be the modulation information for the data to be transmitted, data which may be at rates up to 1Mbit/s. At the minimum, the VCO tuning range must cover the bandwidth required for this data rate, given the type of data coding and form of keying employed. In practice, the range must also be broad enough to center the oscillator in the desired frequency band regardless of processing variations; however the mechanism used for this calibration need not be the same one activated by the input data signal. In the framing of the synthesizer architecture, this signal was assumed to be controlling an on-chip varactor, thus the design is based upon a low gain and allows for wide tolerances. The specified values, for an input control voltage range of 0 to 3V (relative to chip ground), are summarized in Table 4-1.

While the input to the VCO will be a signal at a comparatively low frequency, the output generated will be the RF signal in the 1.8GHz band. In the proposed transmitter architecture, this output signal will need to drive a power amplifier and also a frequency

divider (in the feedback loop), perhaps through some distance of interconnect. On the scale of an integrated circuit, these distances will be electrically short (in silicon, 1 mm is about $\lambda/100$ at 2GHz), so the effects of the interconnect metal will be primarily resistive and capacitive loading, rather than creating signal reflections due to characteristic impedance mismatches. So while buffering will be needed to protect the VCO from capacitive loading, and to isolate it from switching in the divider circuit, the buffers required for the integrated transmitter can be designed for voltage signals. This paradigm will not hold for the initial turn of test silicon, since the capability to probe the RF signal is desired, and also because the synthesizer proof of concept test bed will not be fully integrated. As a result, the LMVCO chip will require impedance-matched output buffers.

Table 4-1: VCO Gain Specifications

<i>SPECIFICATION</i>	<i>VALUE</i>
Input voltage range (V_{in})	0 to 3 V
Nominal VCO gain @ $V_{in}=1.5$ V	20MHz/V
Maximum VCO gain	40MHz/V
Minimum VCO gain	10MHz/V

4.2 VCO Circuit Design

When confronted with a new problem, one of the first things a circuit designer considers are the devices which are available to solve that problem. As briefly covered in the initial chapter, for this project, a silicon bipolar process was chosen that was designed with the goal of allowing for low power, low noise, RF and microwave circuits. While the ADRF process provides a rich array of circuit elements, the components required for a harmonic VCO will draw from a select few of the offerings. To implement the positive feedback mechanism, transformer-like magnetically coupled line sections could be constructed in the metal levels, but lumped capacitors are more efficient at the frequencies under consideration. Metal-insulator-metal (MIM) capacitors would be optimal for this work, but the only structures available are of the poly-n+ variety. These capacitors will also form a portion of the tunable resonator, in conjunction with inductors characterized in the previous chapter and one of several possible p-n junctions for use as a varactor. Per-

haps most importantly, for the active gain elements, the vertical NPN transistors are the only real choice for their high frequency and low noise capabilities, although both NMOS and lateral PNP devices are accessible in ADRF for biasing and switching purposes.

4.2.1 Transistor Considerations

Choosing transistor sizes for use in the gain stage of the oscillator is an exercise in balancing noise performance and power consumption. Bipolar devices with large emitter areas are preferred in order to minimize the ohmic resistances in the base and emitter which contribute significantly to noise figure. However, large transistors need large bias currents to be “fast” enough, both in terms of f_{MAX} and f_T , at the chosen current density. To support oscillation, a transistor must have greater than unity power gain at the desired frequency, which requires sufficient bias current to fulfill the condition of f_{MAX} being greater than the frequency of oscillation. (In quantifying transistor performance, f_{MAX} is also referred to as the maximum oscillation frequency.) In addition, the unity current gain frequency (f_T) of the device has a strong dependence upon bias current density, which in turn sets the transistor current gain at the frequency of interest ($\beta(f) = f_T/f$). Partly as a result of this relationship, the noise figure of a transistor is a function of its collector current, for which there is typically an optimum value. Below this optimum bias level, reduced beta and increased values of the transistor small-signal resistance—through which the collector shot noise current flows—conspire to yield greater noise figures. For higher than optimal bias currents, the increase in noise results largely from increased base (shot) noise current flowing in the source and extrinsic base resistances, but current-crowding and base push-out (the Kirk effect) also contribute. As a result of these noise and speed concerns, the bias current levels in the VCO are constrained; the key to reducing power, then, is clearly by reducing the supply voltage.

4.2.2 Negative Resistance Element Topology

Being able to support a large signal swing with a meager supply voltage is the first topological issue to be addressed. The RF power generated by an oscillator depends on the ability of its gain element to sustain losses in the resonator, feedback structure, and load under large-signal conditions. One way to create an amplifier stage which accommodates a large output signal is by biasing the collector of an NPN through a frequency-

selective element that is low impedance at DC, and high impedance at the desired RF frequency. A quarter-wavelength line (usually combined with radial stubs and lumped capacitors at the supply end) is often used in narrowband microwave systems, but a basic inductor also has this property. Since an inductor cannot support an average (DC) potential, the full supply voltage is impressed upon the collector of the transistor. While this aids low voltage operation and removes one source of power dissipation, taking the RF signal at the collector also maximizes the available room for output swing by centering it around the highest voltage in the circuit. (Swinging above the supply rail does not present junction bias problems for the collector of an NPN built in a p- substrate.) In addition to providing a DC current path, the collector inductor may be designed as part of the resonator—a vignette which will be discussed further. One possibility embodying these concepts is sketched below in Figure 4-1.

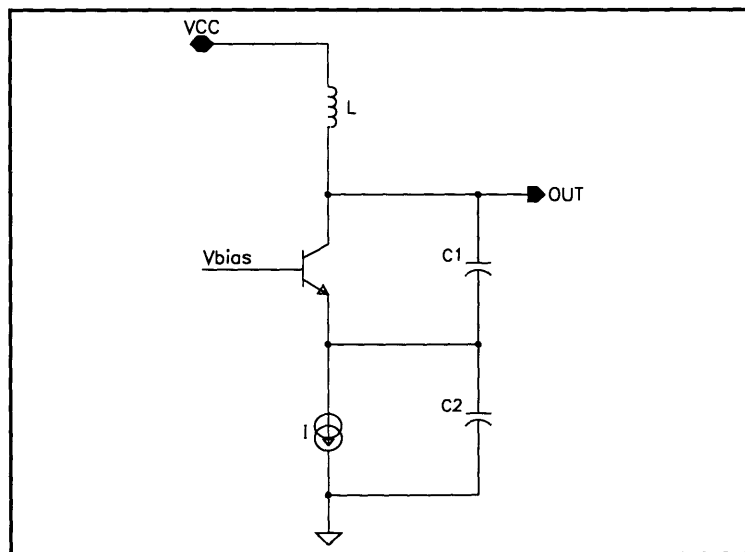


Figure 4-1. Single-ended common-base oscillator circuit.

In this circuit, a common-base amplifier stage provides positive gain from the emitter to the collector, allowing the feedback provided by the capacitive divider to be of the positive sense. These elements comprise a negative resistance generator, to which a resonator may be attached. Initially taking the collector inductor to be large, a first-order, small-signal analysis (with transistor parameters g_m , r_π , and C_π) of the generator yields the impedance looking into the output port to be:

$$\begin{aligned}
Z_{out}(s) &= \left(\frac{1 + sC_{\pi}r_{\pi}}{r_{\pi}} + g_m + s(C_1 + C_2) \right) \left(\frac{1}{sC_1} \right) \left(\frac{r_{\pi}}{1 + sC_{\pi}r_{\pi} + sC_2r_{\pi}} \right) \\
&= \left[r_{\pi} \left(1 + \frac{C_{\pi} + C_2}{C_1} \right) + \left(\frac{1}{sC_1} \right) r_{\pi} \left(g_m + \frac{1}{r_{\pi}} \right) \right] \left(\frac{1}{1 + s(C_{\pi} + C_2)r_{\pi}} \right).
\end{aligned}$$

Separating out and then simplifying the real components renders the condition for a negative resistance at the output port ($\Re \{Z_{out}(s)\} < 0$):

$$\beta > \frac{C_1}{C_{\pi} + C_2},$$

where β is the low frequency transistor current gain. Hence by choosing capacitor values such that $C_2 > C_1$, a negative resistance is ensured for all frequencies where the model remains valid. An oscillator may then be constructed with this element by connecting a resonance to the output port. A stability analysis would show that a parallel resonator is compatible with this negative resistance element, thus the existing devices may be utilized. The resonant frequency for the parallel combination of the collector inductor and the capacitive voltage divider will largely determine the rate at which the circuit will oscillate (in this configuration, C_{π} is in parallel with C_2):

$$2\pi f_0 = \frac{1}{\sqrt{L \frac{C_1 C_2}{C_1 + C_2}}}.$$

Another possible topology which can make use of the collector inductor bias scheme is the same Colpitts-style oscillator, but with an emitter-coupled pair gain stage as pictured in Figure 4-2. Qualitatively, the operation of the negative resistant element in this circuit may be understood by noting that for a voltage source applied to the output node, the current drawn is determined by one transistor in the pair (Q_L), but changes in the source voltage are fed to the base of the opposite transistor (Q_R). As the voltage applied to the output port increases, more of the tail current flows through transistor Q_R , and less from the output node. Thus, a negative resistance is realized, to which a resonator—that again may be formed by the parallel combination of the collector inductor and the capaci-

tive feedback circuit—may be attached to create an oscillator. (As further evidence, the negative resistance of an emitter-coupled pair structure such as the one discussed here was analyzed numerically as an example in Chapter 2.)




Figure 4-2. Single-ended ECP oscillator circuit.

In this form, the differential pair offers two primary benefits over the single transistor common-base variant: less distortion in the output, and the ability to take the output signal as the collector current of Q_R , thereby incurring no loading of the resonator. However, this application will ultimately require a fully differential oscillator, in which form the two circuits are equivalent on these counts. Furthermore, there are a number of important advantages to the common-base implementation. First, all of the nodes in the common-base oscillator are low impedance for low frequencies; this will help keep substrate noise due to adjacent circuitry from modulating the RF output. Along the same lines, the bias scheme is much cleaner—since base bias resistors not required for the common-base gain stages, a lower noise circuit should result. Another benefit of common-base amplifiers over the common-emitter configuration is that base-collector capacitance is grounded at one end, so the high frequency performance will not be hindered by the Miller effect (a problem in fully differential versions of the emitter-coupled pair oscillator). Finally, the noise matching analysis conducted in the following section indicated some favorability with regards to the ease of which the common-base oscillator circuit could be resolved and tuned for low noise performance—a meaningful design issue.

4.2.3 Noise Matching

By deriving the impedance looking into the negative resistance port of the common-base Colpitts oscillator, it was demonstrated that this circuit will create poles in the right half of the s -plane for a broad range of capacitor ratios. It is worthwhile to investigate if this flexibility may be used to buy some improvement in the oscillator performance by noise matching. Providing a noise figure optimizing input match to the gain stage in the oscillator is a linear problem, and may be handled apart from other noise concerns.

For the purposes of small-signal modeling, the feedback action of the capacitors is sometimes replaced by an ideal transformer because of the fixed ratio between the AC voltages at the output and the input. However, this is not accurate in all cases. Namely, the transformer also dictates that if the voltage at the input of the gain stage is v_{out}/n , then the current flowing from the capacitors into the emitter is n times that which flows into C_1 . Rather than rely upon this artifice, a small-signal model was again constructed for the common-base oscillator as done in the previous section, except now with the resonator taken into account. The transfer functions between each of the primary noise sources in the oscillator circuit and the RF output port were then derived. With L and R respectively denoting the values of the inductance and the equivalent resistance in the resonator (refer to Section 2.3), the characteristic polynomial of the system is found to be:

$$D(s) = \{r_{\pi}(C_{\pi} + C_2)LC_1\}s^3 + \{r_{\pi}(C_{\pi} + C_1 + C_2)\frac{L}{R} + LC_1\}s^2 + \{r_{\pi}(C_{\pi} + C_1 + C_2) + (\beta + 1)\frac{L}{R}\}s + (\beta + 1).$$

In this analysis, β is not the high-frequency current gain (i.e. $\beta \neq \beta(f)$), but rather is used as a contraction for $g_m r_{\pi}$. The transfer function between the thermal noise voltage in the extrinsic base resistance ($\overline{v_b^2}/\Delta f = 4kTr_b \text{ V}^2/\text{Hz}$) and the oscillator output port is:

$$H_{v_b}(s) = \frac{v_{out}(s)}{v_b} = \frac{\{r_{\pi}C_{\pi}LC_1\}s^3 + \{L(C_1 - \beta C_2)\}s^2}{D(s)},$$

for the current in the base due to shot noise ($\overline{i_b^2}/\Delta f = 2qI_B A^2/\text{Hz}$):

$$H_{i_b}(s) = \frac{v_{out}(s)}{i_b} = \frac{\{r_\pi LC_1\} s^2 + \{\beta L\} s}{D(s)},$$

for shot noise current in the collector ($\overline{i_c^2}/\Delta f = 2qI_C A^2/\text{Hz}$):

$$H_{i_c}(s) = \frac{v_{out}(s)}{i_c} = \frac{\{r_\pi L (C_\pi + 2C_1 + C_2)\} s^2 + \{(2\beta + 1)L\} s}{D(s)},$$

and for the thermal noise voltage due to the resonator loss ($\overline{v_R^2}/\Delta f = 4kTR V^2/\text{Hz}$):

$$H_{v_r}(s) = \frac{v_{out}(s)}{v_r} = \frac{\{r_\pi (C_\pi + C_1 + C_2) \frac{L}{R}\} s^2 + \{(\beta + 1) \frac{L}{R}\} s}{D(s)}.$$

Assuming that the noise sources are uncorrelated, all of their contributions at the oscillator output port are then summed, yielding a spectrum given by:

$$\overline{v_{out}^2} = H_{v_b}^2(s) \overline{v_b^2} + H_{i_b}^2(s) \overline{i_b^2} + H_{i_c}^2(s) \overline{i_c^2} + H_{v_r}^2(s) \overline{v_R^2}.$$

When plotted as a function of frequency, this equation delineates a power spectrum composed of the additive white noise sources in the circuit, and shaped by the oscillator. This spectrum represents the noise floor of the oscillator at the output port, and it may be integrated to achieve the phase noise (e.g., see Craninckx and Steyaert [1] or Sauvage [2]). However, a more intuitive feel may be gleaned from Leeson's work [3]. As presented in Chapter 2, Leeson predicted a slope of -20dB/decade in the oscillator spectrum around the carrier resulting from a phase-to-frequency translation that arises from the response of the resonator coupled with the feedback mechanism. This characteristic will continue until its line intercepts the circuit noise floor which, Leeson argued, should occur at the half-power corner frequency of the resonator. Thus, in reducing the system noise floor calculated with the expression above, the rest of the oscillator phase noise spectrum should be "pulled" down along with it.




Figure 4-3. Noise floor at oscillator output port with feedback ratio as a parameter.

As a starting point, the base resistance and small-signal parameters were calculated with ADICE from a model for a transistor with a single $0.8\mu\text{m}$ by $10\mu\text{m}$ emitter and biased with a collector current of 1mA . These values, along with an oscillation frequency of 1.8GHz , a resonator with a Q of 10 , and 1mW of output power, were entered into the noise floor analysis, and evaluated with the feedback ratio ($n = V_{out}/V_{in} = 1 + C_2/C_1$) being varied; the results are plotted above in Figure 4-3. Consistent with Lindenmeier's supposition [4], the noise floor is seen to be lowered as the transformer (feedback) ratio is increased from unity to a value of about two, and then begins to rise for larger values of n . This analysis tool is a vital part of the design iteration process which will become more fully developed in the following section. Once the total feedback capacitance is set to obtain the desired frequency of oscillation, the output noise floor is then computed to optimize the ratio of the capacitors. This procedure is embodied by the following design relations:

$$C_2 = nC_{total},$$

$$C_1 = \frac{C_2}{n-1}.$$

4.2.4 Differential Oscillator Design

Previous sections have identified benefits and methods of analysis for a common-base Colpitts oscillator, but a fully differential version needs to be developed due to the high levels of integration which will be realized in the completed product. Merely appending a second copy of the single-ended circuit will not work, as this only results in two freely running oscillators with an indeterminate phase difference. Some form of coupling is needed, and one possibility is a capacitive connection between the emitters as shown below. The addition of this branch allows for one side of the oscillator to draw current at the expense of the other, which is precisely the mechanism needed to achieve the 180° phase split.




Figure 4-4. Differential common-base Colpitts oscillator circuit.

To verify that this circuit does indeed produce a negative resistance between the output nodes, a small-signal analysis was carried out as done earlier for the single-ended design. This time, the condition for the existence of a negative resistance was found to be:

$$\beta > \frac{C_1}{C_\pi + 2C_2},$$

which again covers the frequency span over which the model is valid. For added insight into the effects of the feedback structure, the linear analysis was used to explore the

dependencies of the impedance at the output port upon capacitor ratios and total (series) capacitance. Using the device parameters as calculated for the noise analysis of the preceding section (except with L made large), first the amount of capacitance in the feedback voltage divider was fixed and the transformation ratio was varied. The results for $n \in \{1.5, 2, 2.5, 3\}$ are plotted in Figure 4-5, where it is observed that the negative resistance generated (provided that it exists) depends only weakly upon the capacitor ratio. Next, this ratio was fixed at $n = 2$, and the total divider capacitance was scaled. As shown below, the impacts here are more pronounced—significant even for variations in the capacitance of 20%. The importance of this finding has to do with the selection of the resonator, the relevance of which will become apparent shortly.

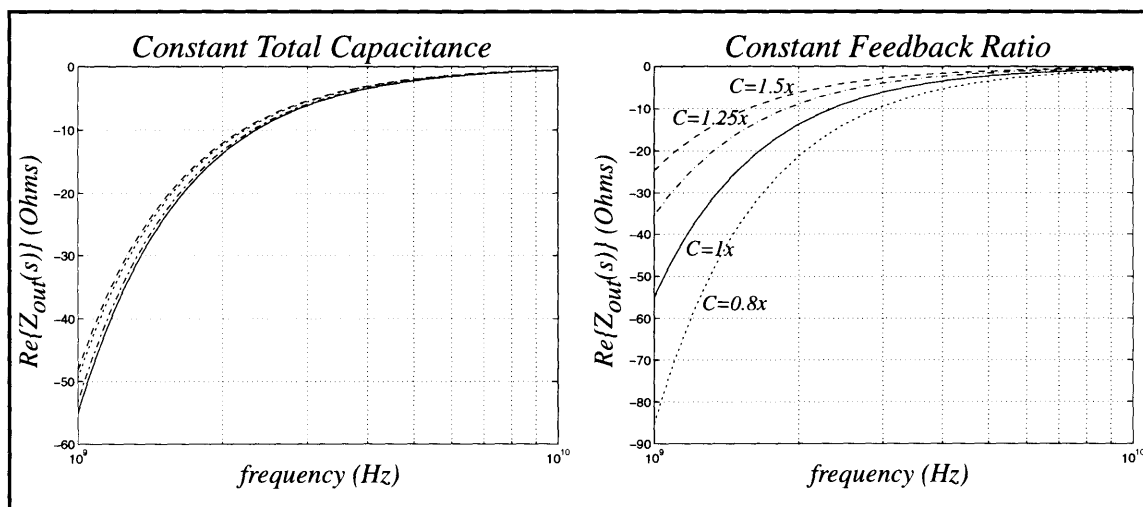


Figure 4-5. Effects of the feedback structure on negative resistance.

It is now expeditious to revisit transistor sizing issues, and to put numbers behind some of the thoughts expressed earlier. The first detail at which to look is that of transistor geometry. When the noise figure is dominated by base resistance (as is generally the case for silicon bipolar devices), it is expected that the lower noise should result by maximizing the emitter perimeter-to-area ratio and the surrounding base contact area. This helps to reduce the series resistance and alleviate current crowding effects, thus transistors are often seen composed of a long emitter stripe of the minimum allowable width. Better performance can be achieved however, by dividing the long stripe into numerous smaller emitters, and then interspersing them with base contacts. To verify this, the input noise

voltage for a common-base gain stage using an ADRF NPN transistor with $8\mu\text{m}^2$ of emitter area was calculated in ADICE over a sweep of bias currents for devices with 1, 2, and 4 emitters of minimum width ($0.8\mu\text{m}$). Since flicker noise is not included in the transistor models, the computed input noise spectrum is uniform over frequency up until the point where device gain begins to fall. One low frequency point from the spectrum of each device was selected at each bias level, and the results plotted below.

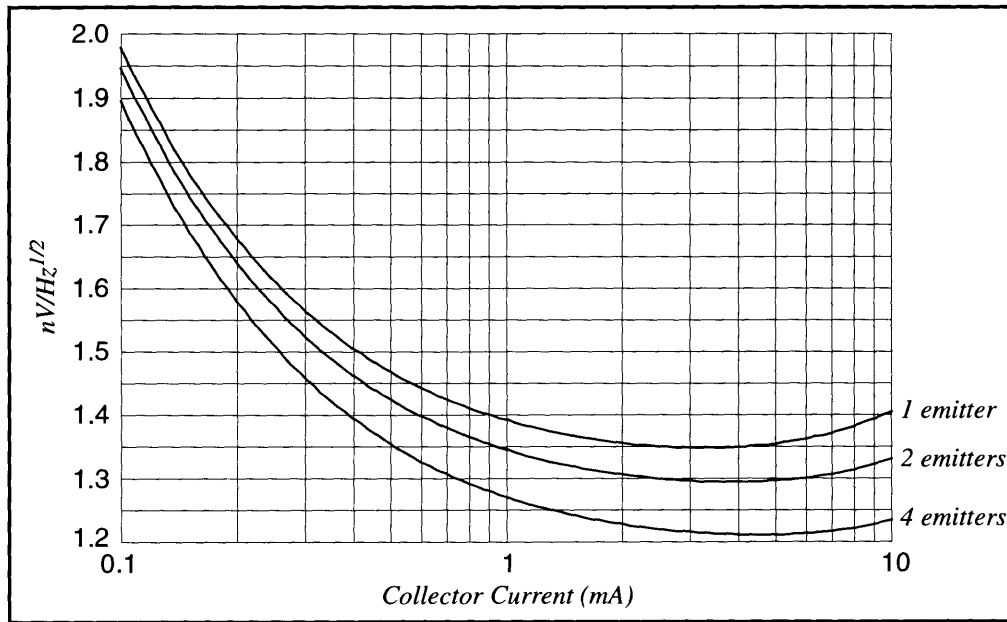


Figure 4-6. Emitter perimeter/area ratio effects on NPN input noise spectral density.

A low frequency point was chosen because it is the transistor’s baseband spectrum which is modulated about the oscillator carrier, however reductions at this point also correspond to a decrease in the minimum noise figure (NF_{MIN}) that may be achieved by the device with an optimum noise match (Γ_{OPT}). As a point of reference, the thermal noise voltage spectral density of a 50Ω resistor at room temperature is $0.9\text{ nV}/\sqrt{\text{Hz}}$. In a 50Ω environment with an input properly matched to the applied signal, the spectral densities of $0.5\text{ nV}/\sqrt{\text{Hz}}$, $1\text{ nV}/\sqrt{\text{Hz}}$, and $1.5\text{ nV}/\sqrt{\text{Hz}}$, translate into noise figures of 1.17 dB , 3.49 dB , and 5.77 dB , respectively.

Although using smaller emitters to constitute a given area does help, the noise levels indicated in Figure 4-6 for $8\mu\text{m}^2$ transistors are still quite high, and process limitations constrain the perimeter-to-area ratio which can be achieved (through the minimum




Figure 4-7. Emitter area effects on NPN input noise spectral density.




Figure 4-8. Emitter area effects on NPN unity current gain frequency.

allowed area for an emitter opening). The only way to further reduce the noise spectrum is by using larger transistors. Keeping the smallest emitter size from the previous simulation, the input noise spectral density for devices with an increasing number of stripes was recorded in Figure 4-7 on the preceding page. The noise density levels continue to be pushed down as more emitters are added (18 in the largest device shown), albeit with diminishing returns—as the amount of noise from the parasitic base resistance decreases, other noise sources become significant. Not surprisingly, however, the reduction of noise via larger devices comes with a penalty. As illustrated in Figure 4-8 (notwithstanding the glitch in modeling), the bigger transistors are slower (smaller f_T) for a given bias current.

Of course, for a low power design, this begs the question of what is fast enough. The first criterion is that the transistor must be “fast enough” to have power gain at the desired frequency in order to support oscillation. As derived in [5] and [6], the unity power gain frequency may be approximated in terms of f_T as:

$$f_{MAX} = \frac{1}{2} \left[f_T \left(\frac{1}{2\pi r_b C_\mu} \right) \right]^{1/2},$$

where r_b is the extrinsic base resistance and C_μ is the base-collector depletion capacitance. For some of the larger devices under consideration, f_{MAX} , as calculated from this expression, has been graphed in Figure 4-9 on the following page. The numbers here are overly optimistic, being about one-third higher than actual, published results for similarly-sized transistors in this process [7]. Even heeding this warning, it must also be kept in mind that this figure of merit is an extrapolated value from measurements at lower frequencies. Some margin, then, should be built into the design to ensure that higher order, deleterious effects have not excessively taxed the power gain at the frequency of desired oscillation. In addition, this factor of safety should also make provision for the losses in the resonator, thereby providing for reliable start-up and amplitude of oscillations.

Similar concessions should also be made with regards to the nominal value of a transistor’s unity current gain frequency. The models and circuits used in most design work with bipolar transistors assume a substantial current gain (i.e. $\beta \gg 1$) in the forward-active regime of operation. Without an adequate margin in the design frequencies, the models, and the design techniques and circuit simulations for which they form the basis,

are rendered inaccurate. Since f_T is lower than f_{MAX} for the ADRF NPNs (as it is with most silicon bipolar transistors), this is usually the more stringent axiom. While it may be possible—through multiple turns of silicon—to achieve a functional oscillator that operates near the unity power gain frequency of the devices within the circuit, where permissible, designs should back away from this edge in order to have a reasonable chance of succeeding.




Figure 4-9. Emitter area effects on NPN unity power gain frequency.

Based on these considerations of noise, speed, and power, as quantified by the data presented, transistors for the oscillator gain stage were chosen to be sized with $15, 2.5\mu\text{m}$ by $0.8\mu\text{m}$ emitters, and nominally biased with 2mA of collector current. This selection should result in a value of around 10 for the f_T over design frequency ratio, and a noise spectral density of $0.7 \text{ nV}/\sqrt{\text{Hz}}$ (corresponding to a noise figure of slightly over 2dB in a 50Ω environment). Accommodations are made, however, in the VCO test chip, to alter the bias on the common-base transistors in the oscillator, thereby allowing some further optimization in this multi-dimensional space.

Now that the gain stage devices have been set, a large-signal analysis of the negative resistance that they generate (when used in conjunction with the capacitors from the previous analyses) may be performed. Again, the collector inductors are made immense in value, and a time-varying, sinusoidal source of 1.8GHz is placed between the differential outputs. Transient simulations are run in ADICE, and the impedance looking into the output port is extracted from a Fourier analysis. Both the impedance and the admittance are plotted below as a function of signal (oscillation) amplitude.




Figure 4-10. Negative resistance of differential Colpitts with increasing signal amplitude.

The most important observation is that *both* the negative resistance and the negative conductance degenerate as the signal builds. As noted in Chapter 2, this will be the case when the magnitude of the imaginary part is larger than that of the real component. In this instance, Kurokawa's stability criterion will be satisfied with either a parallel or a series resonance. This does not imply, however, that both will work equally well. Of considerable importance is the amount of the negative resistance or negative conductance that is produced. In this example, an oscillation may be supported with a series resonator provided that it has less than about 21Ω of resistance. Given a typical, lumped element, resonator with a characteristic impedance of 30Ω , this level of negative resistance will allow for a quality factor as low as 1.43. Using the same inductive and capacitive elements in a parallel configuration, the resonator conductance needs to be smaller than about 4.6mS for an oscillation to begin, which corresponds to a Q of 7.2. Given the similar rates of degra-

duction in resistance and conductance with increasing amplitude, in this form, the negative resistance circuit should support a larger output signal with a series resonance. However, to use the collector inductor as an integral piece of the resonance, the chosen topology is constrained to a parallel resonant circuit. Herein lies the value of being able to change the total capacitance in the feedback structure as demonstrated in Figure 4-5. This mechanism allows the negative resistance element to be more optimally tuned to drive the desired form of resonance. An increase in the feedback capacitance allows for a greater negative conductivity.

With this analysis performed, another design iteration may take place. The transistor sizing and the selection of their bias level in an effort to reduce the achievable noise figure (NF_{MIN}) affected the base resistance and junction capacitances as well as the small-signal device parameters. As a result, the oscillation frequency shifts (particularly with respect to C_{π}), as does the optimum feedback ratio for noise matching purposes. The amount of feedback capacitance is first adjusted to center the oscillator (time-domain simulations are helpful here), and then the capacitor ratio is optimized through the linear noise analysis (Section 4.2.3). As just discussed, these actions change the negative resistance developed by the amplifier, so the effects upon oscillation start-up, amplitude, and stability need to be evaluated. Transient oscillator simulations help to gauge this, but the impedance at the output is also recomputed and then compared with the loss in the resonator. Additional iterations continue from this point until a satisfactory design is reached.

4.2.5 Tunable Resonator Design and VCO Interface Issues

Thus far in this circuit design, a lot of discussion has focused on the capacitive voltage divider which forms the feedback network as well as part of the resonator. Little mention has been made of the resonator's other half, yet the inductors constitute a vital piece of the oscillator. In fact, the value of the inductors was chosen during the design process just detailed, once an idea was acquired regarding the amount of resonator capacitance would be needed to achieve the required tuning range, negative resistance, and quality factor. However, for accuracy in modeling and simulation, the selection was taken from among the discrete units available on the test wafer discussed in Chapter 3. With the choice of the 2mm one-port bond wire inductor (about 2.1nH at 1.8GHz), a lumped model for it was incorporated into the simulations. This leaves the capacitance in the res-

onator as the one remaining degree of freedom for bringing the oscillator frequency into the band of interest.

In the spirit of up-integration, tuning within this frequency band is accomplished by means of an on-chip varactor. The depletion capacitance of a p-n junction under reverse bias is described by a function of the form:

$$C_J = C_{j0} \left(1 + \frac{V_A}{V_{bi}} \right)^{-m},$$

wherein V_A is the applied reverse-biasing potential and C_{j0} is the zero bias junction capacitance which scales with diode area. The remaining parameters—the grading exponent and the built-in potential—are fixed by the doping profile that comprises the junction. To maximize its effect on the resonator center frequency, and to avoid changing the feedback ratio while tuning the oscillator, the varactor is placed in parallel with the bond wire inductors and the capacitive voltage divider. Then, by using the applied voltage as the mechanism by which to tune the frequency, an expression for the oscillator gain may be derived as follows:

$$K_v \equiv \frac{\partial f_{osc}}{\partial V_A} = \frac{1}{2\pi} \left(\frac{\partial}{\partial V_A} (LC_{tot})^{-1/2} \right) = \frac{1}{4\pi} \left(\frac{m}{V_{bi}} \right) L^{-1/2} C_{tot}^{-3/2} C_{j0} \left(1 + \frac{V_A}{V_{bi}} \right)^{-(m+1)}.$$

For some aspects of the design, it is easier to work with the gain written in terms of the frequency of oscillation:

$$K_v = \frac{m}{2} \left(\frac{f_{osc}}{V_{bi}} \right) \left(\frac{C_{j0}}{C_{tot}} \right) \left(1 + \frac{V_A}{V_{bi}} \right)^{-(m+1)}.$$

In either expression, C_{tot} denotes the total capacitance in the resonator, and consists of a “fixed” part—from the feedback capacitors—plus a portion which varies with tuning due to the varactor. Although the value of C_{tot} is now established by the inductor that was chosen, several methods of affecting the gain are indicated here. From a device standpoint, a variety of p-n junctions exist within ADRF, offering the possibility of choosing among several values for the grading factor and built-in potential. Within the circuit

design, both the varactor size (to change C_{j0}) and the levels of reverse bias voltage may be used as variables to help meet the gain specifications (set forth in Section 4.1). Other important factors include the physical limitations of the junction (reverse breakdown voltage), and also the Ohmic resistances in the device—the varactor is a constituent of the resonator, and thus losses in it contribute to overall quality factor degradation.

In considering p-n junctions for use as a varactor, the obvious starting point is with the available bipolar transistors. The base-emitter junction of the NPN is unsuited for frequency tuning duty because the heavy doping results in a low breakdown voltage (in ADRF, 2V is the maximum reverse bias potential which may be sustained by this junction). More robust is the junction between the base and collector of this device, which can safely withstand 5V. While this will suffice, it is still not a good choice as a varactor for a couple of reasons. The first difficulty arises from the ADRF process flow, in which the intrinsic base and collector regions are implanted through the emitter opening, meaning that the base implant cannot be had without the emitter, and as a result, the full base resistance will be present in the diode structure (which would yield a poor Q). Another hindrance in forming a good varactor with this device is the selectively implanted (intrinsic) collector, a step used to increase the collector doping under the intrinsic base, thereby delaying the onset of the Kirk effect to higher current densities. While this implant provides for a higher transistor f_T at bias levels of interest, it leaves behind a very gradually sloping base-collector junction, which translates into a low gain varactor (i.e. low $\frac{\partial C_J}{\partial V_A}$). A better candidate comes from the lateral PNP device, wherein both junctions have the same composition: the extrinsic base handle region from the NPN forms both the emitter and collector terminals of the PNP, and the collector structure of the NPN is the PNP base. This junction is more steeply graded, with $m = 0.3$, and has the added benefit that the anode contact can be made directly above the junction for reduced resistive loss.

At the time of the design, little characterization data existed for the use of this junction as a varactor. A model was available for a minimum size lateral PNP transistor, which included junction capacitance parameters and terminal resistances that were fit to device measurements. In order to obtain any degree of accuracy in the simulation and design work with regards to frequency tuning, the layout of the device from which the model has been constructed was copied exactly with one exception: the emitter “dot” was

removed, leaving only the collector ring to base junction. The remaining diode is imbued with a zero bias capacitance of 39.0fF; this “emitter-less PNP” structure was used as the unit cell from which the needed varactors were constructed. Other relevant parameters for the unit cell are summarized in Table 4-2. For simulation purposes, other needed SPICE model values were also taken from the lateral PNP base-collector junction data.

Table 4-2: Model Parameter Data for PNP-based Varactor Unit Cell

<i>PARAMETER DESCRIPTION</i>	<i>SYMBOL</i>	<i>VALUE</i>
Zero bias capacitance	C_{j0}	39.0fF
Junction grading exponent	m	0.3
Room temperature built-in potential	V_{bi}	1.1 V [†]
Base resistance	r_b	45 Ω
Collector resistance	r_c	5 Ω

[†] Values listed fit a model to capacitance characterization data, and do not represent physical measurements.

With this information, the design procedure is straight-forward. The total resonator capacitance is known, as is the desired oscillator gain (20MHz/V). A nominal applied voltage is chosen, for which the required C_{j0} is solved. This value of V_A is chosen based on device and circuit considerations. In the transmitter of the wireless sensor, the control signal input to the VCO will be from an operational amplifier in the phase comparator with a 3 V supply, so the range of input voltages will be from chip ground up to this supply level (the op-amp can drive its output to within about 300mV of either rail). Since the VCO design is focused on achieving low power through low voltage (a supply of less than 3V), the end of the varactor diode connected to the negative resistance port will be tied to a DC level intermediate to the input voltage range. To avoid forward biasing the varactor, a level shifting stage is needed for the control voltage input. One possibility is to feed the phase comparator output to the base of a lateral PNP used as an emitter-follower (grounded collector) stage, with a number of diode-connected NPN transistors and a resistor up to a larger (e.g., 5V) supply voltage. In ADRF, a $3V_{BE}$ stack (the PNP and two NPNs) yields 2.3V at 20 μ A of bias current; when driving the cathode terminal, this shift is enough to support a positive supply voltage for the VCO (V_{CCV}) of up to about 2.5V—depending on the RF signal swing—for the entire range of inputs. A schematic of this input stage is included in the Appendix.

Squeezing from the other end to constrain the lower bound of V_{CCV} is the on-chip buffer and frequency divider circuitry that interfaces to the VCO. A common-collector stage follows each oscillator output to protect the resonator quality factor from loading, and to isolate the VCO from the switching action of the divider. In addition to providing the inputs to the ECL latches in the divide-by-two block, this buffering stage must also be designed to power a 50Ω load for first silicon empiricism. A differential pair is the basis of each latch, and with the buffer stage that precedes these latches, the common-mode output voltage for the VCO is required to be at least $2V_{BE} + V_{CE_{sat}}$. With this consideration, and after some preliminary simulation work, $1.8V$ was chosen as the nominal value of V_{CCV} . Given a level shift of $2.3V$ in the input stage, the reverse bias voltage applied to the varactor at the center of the input range will be: $V_A = 1.5V + 2.3V - 1.8V = 2V$.

Once the varactor capacitance is solved using this value, the amount of feedback capacitance needed to comprise C_{tot} can be computed. This calculation must also account for device parasitics—primarily C_π and C_μ in the buffering stages of the frequency divider and in the oscillator, and also the collector-to-substrate junction capacitance (C_{JS}) of the oscillator gain stage transistors. Noise and negative resistance concerns are then checked again as detailed in the previous section. For this VCO design, $C_{tot} = 3.741pF$, yielding $C_{j0} = 1.172pF$ for a mid-range gain of $20MHz/V$. Allowing for the other device capacitances in the circuit, $1.789pF = (C_1 C_2) / (C_1 + C_2)$ remains for the feedback network. The optimal transformer ratio for minimizing noise is evaluated to be $n = 2.8$, which is satisfied by $C_1 = 2.783pF$ and $C_2 = 5.01pF$.

To achieve the required varactor size, 30 of the “emitter-less PNP” unit cells must be used. Unfortunately, in this case, capacitance is not the only attribute of a p-n junction. Parasitic resistances (and, at higher microwave bands, contact and metallization inductances) limit the Q of the device, an effect which can impact the quality factor of the resonator. For moderate frequencies, the varactor Q follows:

$$Q = \frac{1}{f_{osc}} \left(\frac{1}{2\pi R_S C_J} \right) = \frac{1}{f_{osc}} \left(\frac{1}{2\pi (r_b + r_c) C_{j0}} \right).$$

The qualifier “moderate” implies frequencies large enough such that the series resistance (R_S) dominates the response rather than the output resistance in parallel with the junction,

but not so high that the previously mentioned inductances come into play. Since the junction capacitance increases with area while the resistance terms scale inversely by the same factor, the varactor quality factor is independent of size (to the first order).

Using the values listed in Table 4-2, the “emitter-less PNP” should exhibit a Q of 45 at 1.8GHz. While this amount of loss is not negligible, it is significantly less than the contribution of the bond wire inductors (which have a Q of about 13). The remaining resonator components are the poly-n+ feedback capacitors, for which the quality factor is a bit uncertain at the time of the design. The bottom plate of this structure is the n+ collector contact diffusion, and has a sheet resistivity of $30\Omega/\square$; this should allow for capacitors having Qs approaching 30 at 1.8GHz. However, the scalable ADICE models provided for these devices are more pessimistic, yielding numbers closer to 10 at these frequencies. With these capacitor models, and including all relevant transistor capacitances, the impedance of the differential resonator—as seen at the negative resistance port—is plotted below for an input control voltage of 1.5V.




Figure 4-11. Impedance of differential resonator with frequency.

From this analysis, the quality factor may be gauged in two ways: from the bandwidth between $|Z| = \frac{\max(|Z|)}{\sqrt{2}}$ points, or from $Q = \left(\frac{\omega_0}{2}\right) \frac{d\phi}{d\omega}$. Either method yields $Q = 8.2$, which seems consistent with the data for each constituent of the resonator. As a worst-case approach, this quality factor, and the scalable ADICE capacitor models from which it was computed, have been used throughout the design process. On LMVCO silicon, each of the devices used in the oscillator have been separated out into individual test sites; these will be characterized so that circuit performance may be better understood.

In an effort to make further, perhaps future, improvements, two experimental varactor structures have been implemented in the VCO. Although these devices were not well characterized at the time of the design, it is anticipated that they will allow for a higher VCO gain for a given varactor capacitance, or, for this design, more feedback (fixed) capacitance for the specified tuning range. The benefits of moving in this direction were noted in the previous section. In addition to the three varactor structures, a fourth LMVCO chip variation was designed: a version with a differential input control signal. Whether the added common-mode noise rejection this configuration offers will be required in the ultimate system is unclear, but it does avail one other potential advantage: the ability to externally “trim” the VCO transfer characteristic by placing an offset between the inputs. Since this adaptation will incorporate an input circuit based upon a differential transistor pair, another possibility may easily be explored—that of having voltage gain in the input stage. If this voltage gain is labeled K_i , then the overall gain of the oscillator is $K_i K_v$, and this product is designed to the specifications of Table 4-1. Again, with $K_i > 1$, a smaller varactor may be used to accomplish the tuning range goals. Another possibility is that the gain of the input stage need not be constant—it may be made non-linear in a fashion that compensates for the varactor characteristic (such as with $V_{out} = (V_{in})^{1/m}$). Although this compensation capability was not exercised on LMVCO silicon, the final version does employ a gain of $K_i = 2$ between the single-ended VCO control voltage and the differential input signal, allowing the nominal gain of 20MHz/V to be reached with 22 of the “emitter-less PNP” unit cells along with lightly recalibrated feedback capacitor values. Schematics showing device sizes for the differential input VCO, and accompanying input stage, are included in the Appendix.

4.2.6 The Differential Colpitts VCO Circuit

Now that the resonator design has been concluded, all that remains incomplete for the VCO is a means of biasing it. Returning to the circuit of Figure 4-4, current sources and a voltage reference are required in addition to the supply rails. In a power conscious context, simple current mirrors are preferred because emitter degeneration resistors and cascode devices dissipate power. A comparison with ideal sources shows that using simple mirrors does not degrade the amplitude of the oscillation, thus this basic scheme is retained in the final circuit as shown in Figure 4-12. A current gain transistor (Q38) is added to supply the base currents, thereby allowing for more accurate measurement of oscillator current consumption. The reference source pin (*EXTBIAS*) is a low impedance node with a significant amount of capacitance attached, hence stability is not anticipated to be a concern. For ease of testing, the bias resistor has been sized to allow the nominal oscillator bias current of 2mA per side to be realized by externally shorting *EXTBIAS* to the VCO supply pin (*VCCv*). However, the physical distinction has been maintained so that flexibility in experimenting with reduced power levels is not sacrificed.

The output of the VCO is followed by on-chip common-collector buffers driving 50Ω coplanar waveguide RF probe pads from which the 1.8GHz signal will be sensed. It was mentioned in the previous section that *VCCv* needs to be at least 1.8V to support these buffers and the subsequent frequency divider, placing a constraint on the ability to operate the VCO at low supply voltages. Although different buffer/divider techniques should be explored in future revisions, a solution which allows evaluation of this work is to raise the lower supply rail (i.e. the emitters of the current source transistors) to a level above chip ground. While oscillation with supplies down to $2V_{CE_{sat}}$ most likely cannot be achieved since the RF performance of bipolar transistors tends to suffer as saturation is approached, numbers around 1V for *VCCv-NEGv* should be reachable.

With this in mind, the *BASE* bias reference is set to divide the supply voltage across the transistors. For the common-base devices, $V_{CE} = VCCv - (BASE - V_{BE})$, leaving a potential of $V_{CE} = (BASE - V_{BE}) - NEGv$ impressed across the current sources. Maintaining at least 0.5V across each transistor with 1.8V for *VCCv* and a V_{BE} of 0.8V requires the *BASE* reference to fulfill the condition: $1.3V < BASE - NEGv < 2.1V - NEGv$. A *BASE* bias voltage in the lower portion of this range will maximize the available room for RF signal swing, although some margin at the input side of gain element is needed because the amplitude of

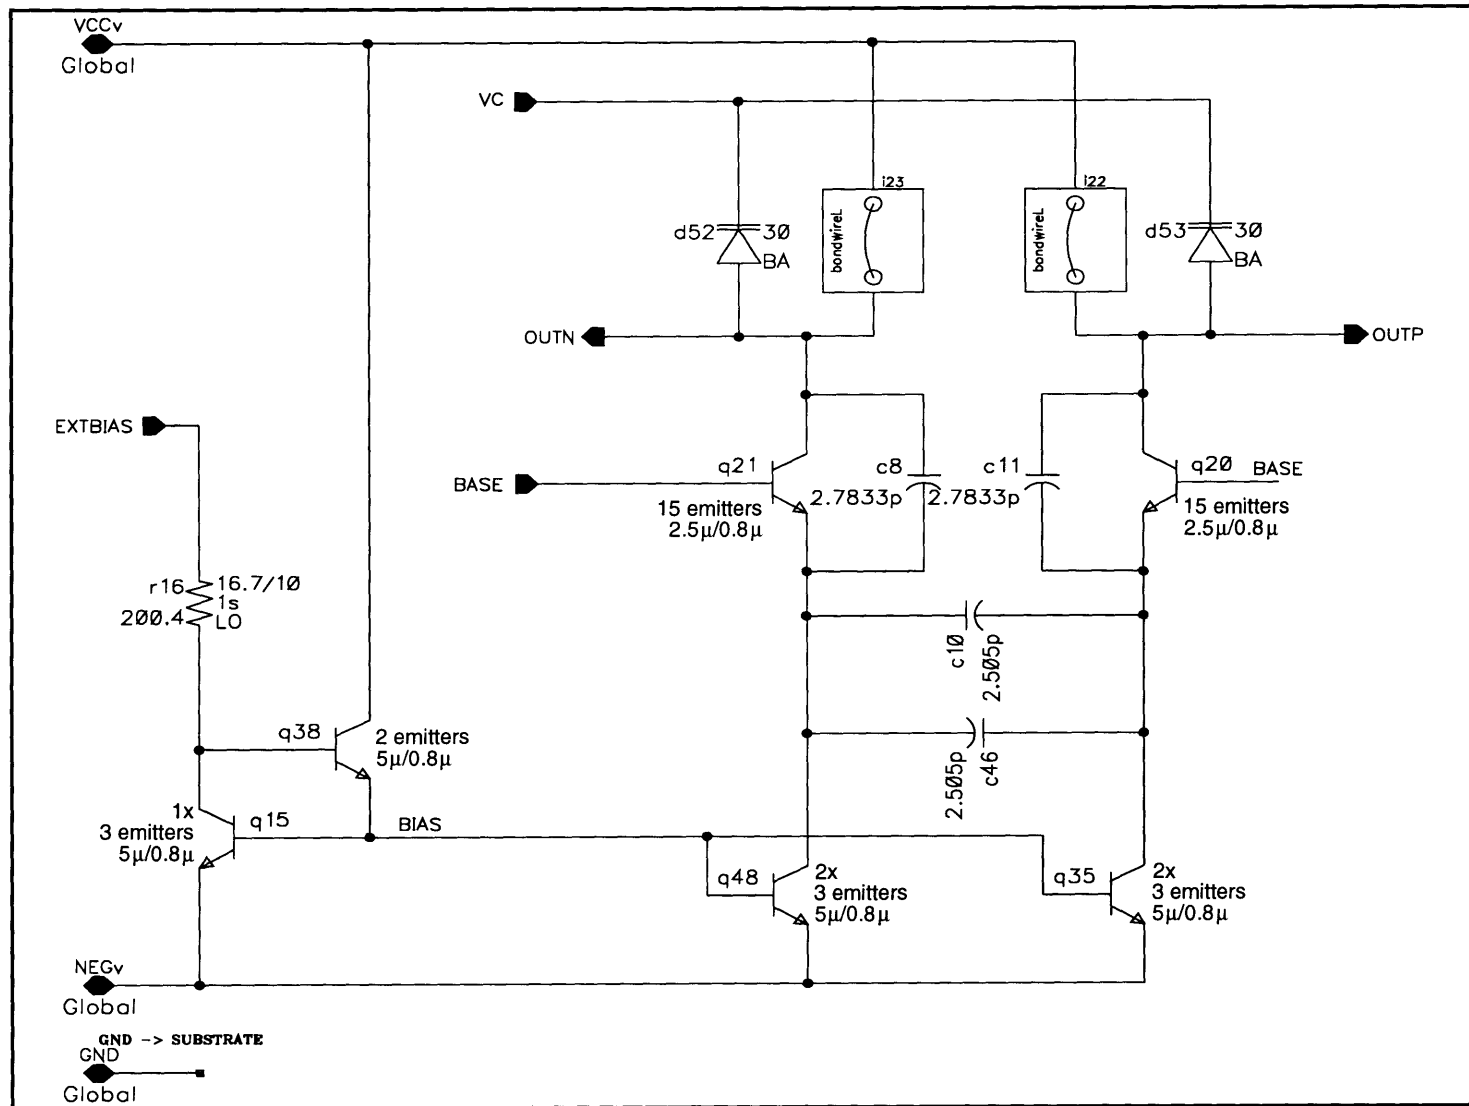


Figure 4-12. Schematic of differential, common-base Colpitts VCO.

the signal there is not insignificant ($V_{in} = V_{out}/n$). For the purposes of design work, 1.5V was chosen as the nominal value of the reference for the common-base transistors, but this node has been routed to a pad in the circuit layout so that experimentation may take place.

A final explanatory note is owed to capacitor C_2 , which is shown in the schematic as two parallel devices. The layout of this capacitor is split to balance the parasitics on either side, and the LVS tool was unable to recognize the end product as a single device. This technique is a good general rule to follow—particularly with the poly-n+ capacitors in ADRF, and helps to ensure that differential circuits remain balanced.

4.3 VCO Simulation Results

Verifying that the VCO circuit reliably builds up and sustains a stable oscillation at a single frequency over a varied operating environment is an integral piece of a robust design. For added pragmatism, all of the simulations presented herein have been conducted with adjacent circuitry in place on both the input and output sides of the VCO: the input level shifting stage drives the varactors, and the emitter-follower buffers and frequency divider are loading the RF output port. The signals plotted in the figures below are taken directly from the VCO output rather than after the buffers, but the primary difference is in the amplitude of the measured signal. Indications of start-up, frequency of oscillation, harmonic content, and relative comparisons are all unaffected by this choice.

The circuit simulation output depicted in Figure 4-13 demonstrates the start-up of the VCO under nominal conditions with typical device models, room temperature operation, a supply of 1.8V, and 2mA of collector current biasing each side of the oscillator. Both output signals are plotted individually in the upper graph, showing the transition to a phase split of 180° as well as oscillation about a common-mode voltage equal to that of the supply. The balanced RF output is traced in the lower plot, reaching steady-state in about 12ns from an initial condition of 1mV. A 4096 point DFT has been used to compute the power spectrum of this signal, is graphed in the figure below the time-domain signals. As shown, the even harmonics are suppressed in the differential signal, while the third and fifth harmonics are -42dBc and -60dBc (decibels relative to the carrier power), respectively. The intent here is to illustrate the relative power in each harmonic; noise in the VCO is not being simulated—the skirts around the carrier and the “floor” of the indicated spectrum are entirely due to the resolution of the FFT.




Figure 4-13. Nominal operation of VCO showing build-up of oscillations.




Figure 4-14. Power spectrum of VCO output signal.

Of course, this is a VCO circuit, so the spectrum should be tunable. For the same nominal conditions used previously, the input voltage to the LMVCO chip is varied, and the frequency of the output is plotted in Figure 4-15. Data was collected from simulations run using the scalable ADICE models for the poly-n+ feedback capacitors. Both single-ended and differential input versions are shown, both of which are tuned to achieve a VCO gain of 20MHz/V for a 1.5V input. Relying upon a smaller varactor, the VCO with a differential input has a greater disparity in gain over the tuning range, and a slightly wider band of operation (44MHz versus 40MHz) for inputs between 0.5V and 2.5V. Both versions, however, fall within the specifications set forth in Table 4-1.

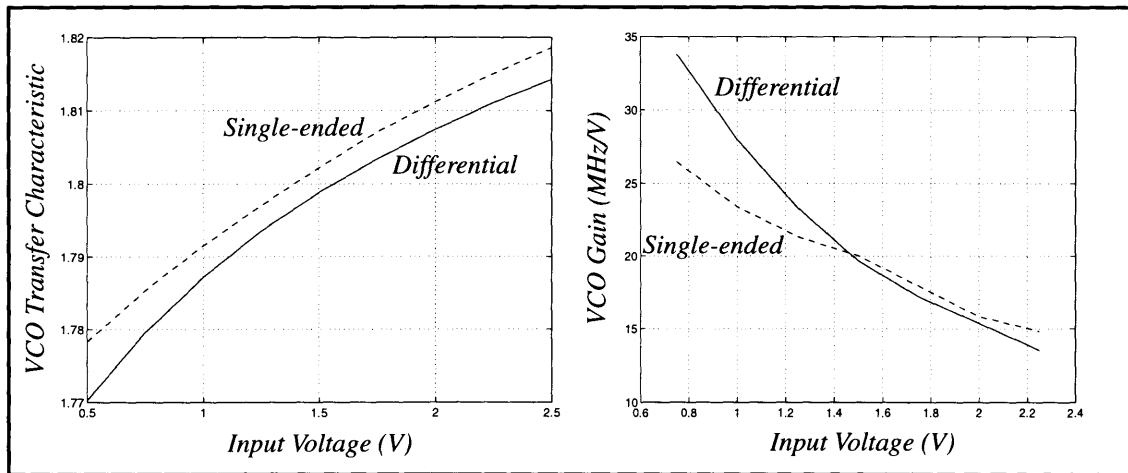


Figure 4-15. Tuning range and gain of VCO.

While this verifies the VCO aspect, of primary importance for this low power oscillator is operation at reduced bias levels. In addition, simulation of the circuit over a range of bias points helps to identify potential circuit problems as suggested in Chapter 2. Using the nominal conditions applied thus far as a baseline, the balanced RF output signal is plotted in Figure 4-16 for successive increments in NEG_V (corresponding to reductions in the VCO supply voltage). Although the start-up transient is noticeably shortened for larger values of NEG_V , the amplitude of the steady-state oscillation is relatively unaffected until the supply is reduced below 1V; this observation is quantified in Table 4-3. The numbers listed in the amplitude column are for the time-domain signal and not just the fundamental component, but the total harmonic content is low so the two values should be

similar. Since the noise figure of bipolar transistors is generally only weakly dependent upon the voltage across them, the consistency in oscillation amplitude should lead to a situation where the degradation occurring in the phase noise spectrum as the supply is lowered is outpaced by the power savings. However, given an available resonator quality factor, overall system noise requirements will still dictate the lower bound on power.




Figure 4-16. Operation of VCO over a range of supply voltages with $V_{CCv} = 1.8\text{V}$.

Table 4-3: Simulation Data for VCO Output Over Supply Voltage

<i>SUPPLY VOLTAGE</i>	<i>FREQUENCY</i>	<i>SIGNAL AMPLITUDE</i>
1.8V	1.800GHz	1.11V
1.5V	1.797GHz	1.10V
1.2V	1.791GHz	1.05V
0.9V	1.763GHz	0.75V

The simulations detailed on the previous page were conducted with a constant voltage applied to *BASE*, referenced to the lower supply rail (*NEG_v*), set to 1.3V so that operation could be sustained with having just 0.9V across the oscillator. Since *VCC_v* remained at 1.8V (with respect to chip ground) throughout the exercise, the potential seen by the base-collector junctions of the common-base transistors changed as the supply was lowered. This mechanism is responsible for the frequency shifts recorded in the previous table.

As just observed, the oscillator is not strongly affected by the supply voltage. However, the bias current through the circuit yields a different story, as told in Figure 4-17. The large-signal negative conductance characteristic is enhanced by increased current such that it does not degrade as quickly while the RF signal builds in amplitude. This trait is reflected by the magnitude of oscillation for a given level of bias current. Upon start-up, the amount of negative conductance produced by the active element initially “begins” at some small-signal level, and then decreases as the effective transconductance wanes due to large-signal conditions. Increased bias levels simply push out the point at which large-signal effects become noticeable, allowing for larger amplitudes of oscillation before the negative conductance has been reduced to the point where it matches the loss in the resonator. The circuit simulations of varying the collector current are summarized in Table 4-4, with 2mA per side of the oscillator as the baseline. Frequency deviations are again owed to the gain transistors of the VCO, but this time it is the diffusion capacitances that are responsible. These capacitors represent base charge storage, and thus have a value proportional to collector bias current.

The operation of the VCO has proven to be robust with nominal devices over a space of operating points, but it is also important to have a circuit that is tolerant of processing variations. To help deal with this line of inquiry, statistical process corner models




Figure 4-17. Operation of VCO over a range of bias currents.

Table 4-4: Simulation Data for VCO Output Over Bias Current

<i>BIAS CURRENT</i>	<i>FREQUENCY</i>	<i>SIGNAL AMPLITUDE</i>
-50%	1.801 GHz	0.41 V
-20%	1.803 GHz	0.85 V
2mA/side	1.800 GHz	1.11 V
+20%	1.797 GHz	1.35 V
+100%	1.786 GHz	2.18 V

are available for use in ADICE. In addition to the typical devices (TYP), element parameters are skewed to represent the extremes of acceptable wafer fabrication. A best-case speed process (BCS) involves low values for resistors and capacitors, along with fast transistors comprised of reduced junction capacitances and forward transit time, plus increased β and current handling capabilities. The worst-case speed (WCS) device set is selected from the opposite end of the spectrum, while the two remaining model libraries consist of all relevant parameters being on the low edge, and on the high edge, of their statistical ranges, respectively. Although it may be argued that such statistical “corners” are not physically realizable within a process flow, they nonetheless provide a good test for the oscillator.




Figure 4-18. Operation of VCO over process.

A 5 ns window of the steady-state condition with each model set is plotted for comparison in Figure 4-18 on the previous page. While the frequency and amplitude do vary with process, the results are again evident that a single, stable oscillation is produced over the range of tolerances experienced by ADRF. Holding the tuning range data presented in Figure 4-15 up against Table 4-5, it is apparent that either a greater VCO gain or another tuning mechanism will be required in a production part. With the varactor employed in this design, the oscillation is tunable by perhaps ± 30 MHz around a center frequency, which is not enough to compensate for process variations—even given an inductor non-varying in value as assumed in these simulations. Increasing the oscillator gain may prove difficult within the system context because it places stricter requirements on the phase comparator and loop filter. Other possibilities include a “course-tune” control—perhaps another, larger, varactor—driven by a DC signal apart from the phase-locked loop, capacitor trimming, or the method of fitting bond wire inductor lengths to wafer capacitance suggested in the previous chapter.

Table 4-5: Simulation Data for VCO Output Over Process

<i>PROCESS POINT</i>	<i>FREQUENCY</i>	<i>SIGNAL AMPLITUDE</i>
WCS	1.679GHz	0.89V
HIGH	1.700GHz	0.95V
TYP	1.800GHz	1.11V
LOW	1.915GHz	1.23V
BCS	1.950GHz	1.33V

It is plainly seen that the available precision of component values in integrated processes presents a problem for lumped LC oscillators such as this one. For many other monolithic circuits, another measure of accuracy is also critical: the ability to match devices. However, matching is not a significant issue for this circuit because the coupling of the emitters forces the sides to be 180° out of phase, at a common frequency determined by the entire resonant structure. Imbalances on one side of the oscillator can lead to an amplitude error between the halves of the differential signal, but not a phase error, so the balanced signal remains strong and contains just one fundamental frequency component. Mismatches among the feedback capacitors will affect the noise match to the

active element, but holding a capacitance ratio within a few percent is easily accomplished by modern processes. An error of this magnitude will hardly be palpable in the phase noise, causing a degradation of perhaps a few tenths of dB (at most) across the spectrum.

Now that the circuit has been repeatedly shown to oscillate, the other aspect of its performance which needs to be quantified is the phase noise spectrum. The information needed to compute the spectrum using Lindenmeier's derivation has now been resolved for this design. With a center frequency of 1.8GHz, a resonator Q of 8.2 as calculated in Section 4.2.5, and a noise figure of 2dB (Section 4.2.4), the noise spectral density relative to three different carrier power levels is shown below in Figure 4-19. Given the RF voltage signals from the circuit simulations presented throughout this section, an output power of 1 mW (0dBm) into a 50Ω load seems feasible. At an offset of 1 MHz from the carrier, these factors result in a noise spectral density of -125 dBc/Hz.

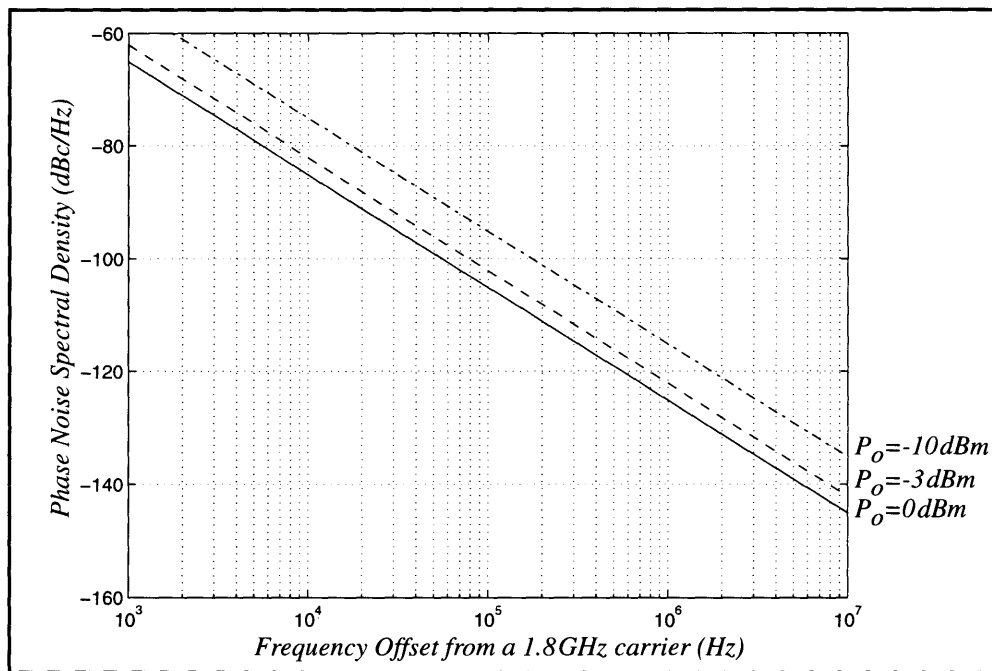


Figure 4-19. Calculated phase noise spectral density using Lindenmeier.

4.4 Frequency Drift

Once achievable specifications for a circuit element can be penned, ideas regarding its use can take form. Consideration of other possible architectures for extremely low power transmitters has led to the thought of operating the VCO open loop, or perhaps

having it controlled by some remote feedback mechanism. Such a conceptual path leads to the question of frequency drift: left to its own devices, what will happen with the VCO output frequency over time? The time scale generally considered as “drift” is an intermediate one: effects included in drift have a longer time constant than would be captured in phase noise analyses and measurements, but long term product degradation (e.g., hot carrier effects, metal migration) is explicitly excluded.

For the purposes of this discussion, the VCO is assumed to have been packaged and nominally tuned so that the frequency is centered. This leaves three run-time variables that can potentially influence the operation of the VCO: supply voltage, bias current, and temperature. As previously demonstrated, the primary result of changes in the bias current is to moderate the signal amplitude along with the noise spectrum, but not the frequency of oscillation which is only weakly affected. From the simulation results printed in Table 4-4, the sensitivity of the oscillator frequency to the bias level is calculated to be 83ppm for a 1% change in current. If large swings in the bias current were present, this effect could be significant, but for typical deviations, this sensitivity term will often be outweighed by other factors. If the fluctuations in current are happening at a faster (than drift) rate, AM noise will result and will augment the phase noise spectrum.

Through the collector inductor, the supply voltage (V_{CCV}) appears as the DC level on one side of the varactor, with the control voltage on the other. Since the capacitance of a varactor is sensitive to the potential impressed across it, a change in V_{CCV} will tune the VCO in the same manner as modulating the input. This is an issue that needs to be addressed through proper system design in the transmitter: the control voltage generated by the phase-locked loop must be referenced to supply. Changes in this voltage then will not appear across the varactor, and thus the oscillator frequency will not be affected. Similarly, if the reference voltage for the common-base transistors in the oscillator moves in step with V_{CCV} , the frequency shifts noted earlier in Table 4-3 can be avoided.

The issue of temperature is more complex. In order to think about this, it is helpful to return to the theory of determining the rate of oscillation: the frequency for which all of the reactances at the negative resistance port sum to zero. Each element that is incorporated as a part of the resonant circuit needs to be examined over temperature. The most prevalent components are the bond wire inductors; these may droop as they warm, but

their reactance is (to the first-order) dependent only upon the bond pads and the length of wire—the wires must be really close to substrate for their height to have an effect. Thus, once in place, the inductors should not contribute to frequency drift.

The first capacitance that inevitably comes to mind is the varactor. Earlier work in deriving VCO gain expressions began with a junction capacitance formulation in terms of C_{j0} that made the applied voltage dependence explicit. This is tedious as a starting point for a drift discussion because C_{j0} has an implicit temperature dependence. It is more fortuitous to return to the parallel plate definition of areal capacitance, from which Sze writes for a linearly graded junction [8]:

$$C_J \equiv \frac{\epsilon_{Si}}{W} = \left[\frac{qa\epsilon_{Si}^2}{12(V_{bi} + V_A)} \right]^{1/3} .$$

As before, V_A is the applied reverse bias potential, a is the slope of the doping profile (in cm^{-4}), and ϵ_{Si} is the permittivity of silicon. Taking the partial derivative with respect to temperature and generalizing yields:

$$\frac{\partial C_J}{\partial T} = -mC_J(V_{bi} + V_A)^{-1} \frac{\partial V_{bi}}{\partial T} ,$$

where m again represents the junction grading exponent. The built-in potential has a complex relationship with temperature through both the thermal voltage (kT/q) and the intrinsic carrier concentration of silicon; evaluating its differential as contained in the prior expression gives:

$$\frac{\partial V_{bi}}{\partial T} = - \left(\frac{2kT}{q} \left(\frac{1}{n_i} \right) \frac{\partial n_i}{\partial T} - V_{bi} T^{-1} \right) ,$$

which involves the sensitivity of the intrinsic carrier concentration (n_i) to temperature. Using a numerical differentiation technique to assess this term at $T = 300K$ provides:

$$\left[\left(\frac{1}{n_i} \right) \frac{\partial n_i}{\partial T} \right]_{T=300K} = 89800 \text{ ppm}/(^{\circ}C) .$$

These expressions may then be solved for the desired temperature sensitivity of the varactor capacitance, which is also a function of the applied potential. Using the parameters in Table 4-2 for the “emitter-less PNP” varactor unit cell, with a mid-range input of 1.5 V (i.e. $V_A = 2V$),

$$\left[\left(\frac{1}{C_J} \right) \frac{\partial C_J}{\partial T} \right]_{T=300K} = 97 \text{ ppm}/(^{\circ}C) ,$$

although this number will increase for lower control voltages ($274 \text{ ppm}/(^{\circ}C)$ with $V_A = 0V$), and decrease for those higher than 1.5 V. This analysis applies not only to the varactor, but also to C_{μ} and C_{JS} (the collector-to-substrate junction capacitance) of the oscillator gain transistors, and to C_{μ} of the buffering devices. All of these capacitances are effectively paralleled with the varactor (although each has a different voltage impressed across it), the sum of which contributes to drift. Of the total capacitance in the resonator, 30% arises from p-n junctions, and is subject to this sensitivity.

The diffusion capacitances within the resonator are also temperature sensitive. From [8], these capacitances resulting from base charge may be expressed (using \Im to denote the imaginary part) as:

$$C_D = \frac{qI_C}{kT} \left(\frac{1}{\omega} \right) \Im \{ \sqrt{1 + j\omega\tau_F} \} ,$$

where τ_F is the forward transit time, which adds a temperature dependence (through the minority carrier diffusivity in the base) to the T^{-1} term in the diode conductance. The frequency dependence indicated in the previous expression is not significant for operation well below the transistor f_T , and may safely be neglected. Remaining, however, is a temperature sensitivity of the diffusion capacitance that is not insignificant when compared with that of the junctions, but the overall effect it has on the oscillator frequency will be smaller, since C_{π} terms contribute only 10% of the total resonator capacitance. Hence, the junction temperature sensitivity may be used to obtain “ball park” numbers for temperature drift, which then may be backed up with simulation results.

The remaining capacitors in the oscillator, the poly-n+ feedback elements and the metal to substrate parasitics, typically have smaller variations with temperature. Although

this aspect of these devices is not characterized for ADRF, nor modeled in ADICE, sensitivity values are usually below $50 \text{ ppm}/^\circ\text{C}$ for MOS capacitor structures, and even less than this for metallization capacitances. Again, this suggests that the p-n junction capacitance sensitivity and circuit simulations should respectively yield reasonable zeroth-order and first-order predictions of VCO response over temperature. Results from simulations at several input voltages are listed in the following table.

Table 4-6: Simulation Results of VCO Frequency over Temperature

<i>INPUT VOLTAGE</i>	<i>OPERATING TEMPERATURE</i>	<i>RELATIVE ΔF</i>	<i>ABSOLUTE ΔF</i>
0.5V	-40°C	+0.8%	+13.8MHz
	30°C	—	—
	125°C	-1.1%	-18.8MHz
1.5V	-40°C	+0.6%	+10.4MHz
	30°C	—	—
	125°C	-0.7%	-12.7MHz
2.5V	-40°C	+0.5%	+9.6MHz
	30°C	—	—
	125°C	-0.6%	-10.0MHz

The amount of drift evidenced by the numbers above suggests that an integrated oscillator, such as the one presented in this chapter, is incapable of being used as an open loop transmitter for most applications. However, the oscillator should not drift so far that intermittent (i.e. when transmitting) or remote feedback mechanisms are out of the question. Caution should be exercised when building systems around VCOs with higher gain, since the drift experienced in the oscillation frequency will also be appreciably higher. Short of employing alternate tuning mechanisms, no methods of significantly reducing drift have been unearthed by this investigation. It has been surmised that the inductance value should not be very susceptible to temperature, thus increasing its proportion of the required LC product may yield some improvement. However, this would come at the expense of gain and possibly resonator Q—neither of which are desirable consequences.

References

1. J. Craninckx and M. Steyaert, "Low-Noise Voltage-Controlled Oscillators Using Enhanced LC-Tanks," *IEEE Transactions on Circuits and Systems—II: Analog and Digital Signal Processing*, December 1995, pp. 794-804.
2. G. Sauvage, "Phase Noise in Oscillators: A Mathematical Analysis of Leeson's Model," *IEEE Transactions on Instruments and Measurement*, December 1977, pp. 408-410.
3. D. B. Leeson, "A Simple Model of Feedback Oscillator Noise Spectrum," *Proceedings of the IEEE*, February 1966, pp. 329-330.
4. H. K. Lindenmeier, "Noise Matching Techniques in Transistor Oscillators," 1980 *IEEE Circuits and Systems Symposium*, pp. 1052-1055.
5. I. J. Bahl and P. Bhartia, *Microwave Solid State Circuit Design*, John Wiley & Sons, New York, 1988, Chapter 7.
6. S. M. Sze, *Physics of Semiconductor Devices*, second edition, John Wiley & Sons, New York, 1981, Chapter 3.
7. K. O, et. al., "A Low Cost and Low Power Silicon *n*pn Bipolar Process with NMOS Transistors (ADRF) for RF and Microwave Applications," *IEEE Transactions on Electron Devices*, October 1995, pp. 1831-1840.
8. S. M. Sze, *Physics of Semiconductor Devices*, second edition, John Wiley & Sons, New York, 1981, Chapter 2.

V. LMVCO Design Recapitulation and Conclusions

A low power, wireless transmitter has been introduced as a backdrop which requires an integrated, low noise, voltage-controlled oscillator (VCO). Toward achieving operation at low power levels while minimizing the output spectrum due to phase noise, resonator based oscillator architectures offer the best performance at microwave frequencies along with minimal active element requirements; the challenge, however, proves to be implementing a monolithic resonator of sufficient quality. With present manufacturing technologies, lumped element LC resonant tanks remain the most viable option for low power oscillation in the desired 1.8GHz band.

Insight into the importance of the resonant element is gained by investigating oscillator models and performance dependencies. The negative resistance port model of an oscillator is generally the least cumbersome design approach to use at microwave frequencies, for which an important distinction is drawn between “voltage-controlled negative resistance” elements (VCNRs) and those which are “current-controlled” (ICNRs). Although this is an important issue, available literature is often unclear about how to make a determination in practical circuit design situations. Intuitive methods are valuable, but a numerical analysis may be performed by using a source of the desired oscillation frequency to drive the negative resistance port. Looking into this port, if the magnitude of the real part of the conductance ($|G|$) about the oscillation frequency is seen to decrease monotonically with signal amplitude, a VCNR element is indicated. Reciprocally, a monotonic decrease in the magnitude of the impedance ($|R|$) is consistent with the ICNR definition. There are, however, cases where *both* $|R|$ and $|G|$ can decrease; the key here is monotonicity over the range of signal amplitudes, the significance of which lies in the stability of the oscillator.

In the condition of steady-state oscillation, a pair of complex conjugate circuit poles will have moved to the $j\omega$ -axis at the frequency where the imaginary parts of the impedances at the negative resistance port sum to zero. However, for this to be a stable operating point, the slope of the phase response of the resonator must be chosen in accor-

dance with the configuration of the negative resistance element. To satisfy this requirement, a series resonance must be used when a device exhibits a monotonic decrease in negative resistance as the signal builds, while an element which monotonically degrades in its negative conductance needs to be mated with a parallel resonator.

While this strategy helps to ensure the desired oscillation, minimizing its noise is another issue. The spectrum produced by an oscillator is dominated by phase noise for all frequency offsets from the carrier of interest. Originally through the work of Leeson [1], it is seen that the quality factor (Q) of the resonator plays a large role in determining noise performance, since the phase noise spectrum is reduced by Q^2 . Power, however, is also a contributor, entering the spectrum both through the noise figure of the oscillator gain stage, and also through the output power in the carrier (P_o) and its association to the DC power consumed in the VCO circuit. Although harmonic oscillators are bound by this inherent coupling between noise and power, the relationship is subject to a number of device and circuit design issues.

One such issue was recognized by Lindenmeier [2], that of providing an optimum noise match to the oscillator gain stage. Similar to the design of RF/microwave low noise amplifiers where a matching network is used to transform the given source impedance to that which minimizes the noise figure, the (positive) feedback ratio (i.e. V_{out}/V_{in}) imposed about the gain stage may be modified to obtain the optimum source impedance (Γ_{OPT}). The ability to minimize phase noise in this fashion is consistent with the concepts of noise matching, and is a valuable tool for improving oscillator performance.

Carrier power and oscillator noise figure have significant effects on the phase noise spectrum, but for low noise, low power applications, it is vital that the relationship of the spectrum to Q^2 be exploited. To help achieve the desired inductor and circuit performance, a set of monolithic structures was designed and fabricated at Microsystems Technology Laboratories using a top-level metal in a typical silicon process flow. Various planar spirals were fabricated and characterized, as were some bond wire inductors constructed from wires which jump from one bond pad to another on the same integrated circuit die. A number of observations from this experiment have proven useful in the design of monolithic inductors.

For the planar spirals, the inductance was seen to be fairly constant for a given number of turns in the spiral and pitch of the turns, but the quality factor is improved as

the conductor width increases and the metal spacing decreases. Another useful detail highlighted experimentally is that an octagonal structure exhibited an improved Q when compared to its square brethren of the same number of turns. The inductance was also diminished (by about 15%), but the increase in Q (20% to 30%) overshadowed this drop. In general, the more circular the structure, the greater the quality factor will be for a given inductance, as resistive corners are alleviated and finally eliminated as a circle is approached. For the same reason, the middle area of the spirals should be left open, since tight turns introduce many corners without adding much inductance. Similar conclusions were drawn regarding the inductor self-resonant frequency (f_{SR}), which tracks Q in many regards. Additionally, it was noticed that f_{SR} for the spirals was more closely linked to the number of turns in the spiral, rather than the total area of metallization occupied by it. Again, this suggests that spiral inductors may be improved by increasing the metal width to spacing ratio.

Overall, the bond wire inductors performed better, yielding quality factors about three times that of the spirals (~13 versus ~4). This improvement is owed to the bond wires, which are both less resistive and farther from the lossy silicon substrate than the metal comprising the spirals. With these advantages, the bond wire structures produced self-resonant frequencies above 20GHz until the longest wire lengths (3mm) were reached. Gold wires offered no discernible advantage over aluminum, suggesting that the primary loss mechanism in the bond wire inductors is through capacitive coupling into the substrate. Both ends of each aluminum wire were attached via wedge bonds, while one bond for each of the gold wires was a ball bond. Thus the results seem to indicate that the smaller footprint area of the wedge overcomes any disadvantages which may accrue due to the non-perpendicular fashion in which this bond attaches the wire to the die. In this study of inductance elements constructed with a manual bonding machine, a sample of 28 inductors with 1.5 mm, 2 mm, and 2.5 mm wires yielded 3σ tolerances on the inductance of $\pm 12\%$. This is comparable to the absolute accuracy achievable for monolithic capacitors, and these numbers could potentially improve with automated equipment adapted for this operation.

Other possibilities include thicker dielectrics to place the inductive elements further above the substrate, reducing the coupling of RF energy into it and thereby increasing

Q. Again, this benefit is provided to any metal structure built upon the deeper insulator, but the designer is limited in this regard by the available process. Thicker metal helps the spiral inductors, but moving to larger diameters will not likely improve the bond wire elements because the skin depth (about $1\ \mu\text{m}$ in aluminum at 2GHz) is a small fraction of even the narrowest wires. Moving to a more conductive medium for the bond wire inductors also brings with it the issue of skin depth, which is inversely proportional to the square root of conductivity ($\delta \propto 1/\sqrt{\sigma}$). Hence, in exchanging aluminum wire for gold, the reduction in the series resistance term will be at most 10%. As long as most of the loss remains in the substrate, improvement through better wires will not be significant.

A bond wire inductor was selected as the cornerstone, and then a VCO circuit was designed around it. The key to reducing power in the circuit is by reducing the supply voltage since the bias current levels are constrained as a result of noise and speed concerns. Using an inductor to furnish the collector current to an NPN transistor helps to create an amplifier stage for the oscillator that accommodates a large output signal with a meager supply voltage. A differential oscillator topology was constructed using common-base transistors, which offer the benefits of not being hindered by the Miller effect at high frequencies and providing low impedances at each oscillator node for low frequency substrate noise. Positive feedback around the gain stages is implemented by a capacitive voltage divider which forms a part of the resonator (along with the collector inductors). The ratio of the capacitors is set to provide a noise figure optimizing input match to the gain stage. Analyzing the noise match for the oscillator is a linear problem, and may be handled apart from other noise concerns.

The transistors for the VCO gain stages must be sized in conjunction with the linear circuit noise analysis. When the noise figure of a device is dominated by base resistance (as is generally the case for silicon bipolar transistors), lower noise results when the emitter perimeter-to-area ratio and the surrounding base contact area are maximized. Although a long emitter stripe of minimum allowable width is often used, better performance can be achieved by dividing the long stripe into numerous smaller emitters, and then interspersing them with base contacts. This helps to reduce the series resistance and alleviate current crowding effects, and thereby the noise figure that may be obtained with an optimal match. Larger devices are less noisy, but also function more slowly at a given

bias. To ensure proper operation, active devices must be designed with an adequate margin in both f_T and f_{MAX} (above the desired oscillation frequency) at an acceptable level of power consumption. This requirement sets an upper limit on transistor sizes, and ultimately on the VCO noise performance.

Before these limits on noise become relevant, the oscillation must be stable. For the differential, common-base, Colpitts oscillator circuit developed for this work, non-linear analyses indicate that both the negative resistance and the negative conductance degenerate monotonically as the signal builds. In this instance, Kurokawa's stability criterion [3] will be satisfied with either a parallel or a series resonance. This does not imply, however, that both will work equally well. Since the chosen low voltage topology with the collector inductors is constrained to the use of a parallel resonant circuit, it is important to adjust the active element for increased negative conductance (as opposed to negative resistance). An effective mechanism for this adjustment has been shown to be the total amount of capacitance in the feedback structure. Increasing this capacitance yields a greater negative conductivity, which should provide for a larger oscillation amplitude.

Of course, the feedback capacitors are included in the resonator, so they also play in setting the center frequency of the oscillation. For a VCO, this frequency must be tunable, and on-chip varactors have been incorporated for this purpose. More steeply doped p-n junctions provide for a greater VCO gain, as do more lightly-doped varactors because of their smaller built-in potentials. Within the circuit design, both the varactor size (to change the amount of capacitance which is voltage-controlled) and the nominal reverse bias voltage may be used as variables to help meet tuning range requirements. Other important factors for consideration include the physical limitations of the junction (reverse breakdown voltage), and also the Ohmic resistances in the device, which contribute to overall quality factor degradation.

A total of four versions of the VCO designed in this work have been implemented in a silicon bipolar technology. This oscillator is being used to explore issues for low power wireless transmitters on a number of levels. To allow for experimentation in minimizing power consumption, each circuit has been designed with the flexibility of individually setting bias levels by an external source. One of the circuit configurations employs a differential control voltage input which offers possibilities of increased common-mode

noise rejection and external trimming of the VCO transfer characteristic, both of which may prove useful in a system context. This version also incorporates voltage gain prior to the oscillator, allowing VCO gain targets to be hit with smaller varactors. The remaining variants are being used to evaluate the performance of different varactor structures against the goals of reduced loss and increased tuning range.

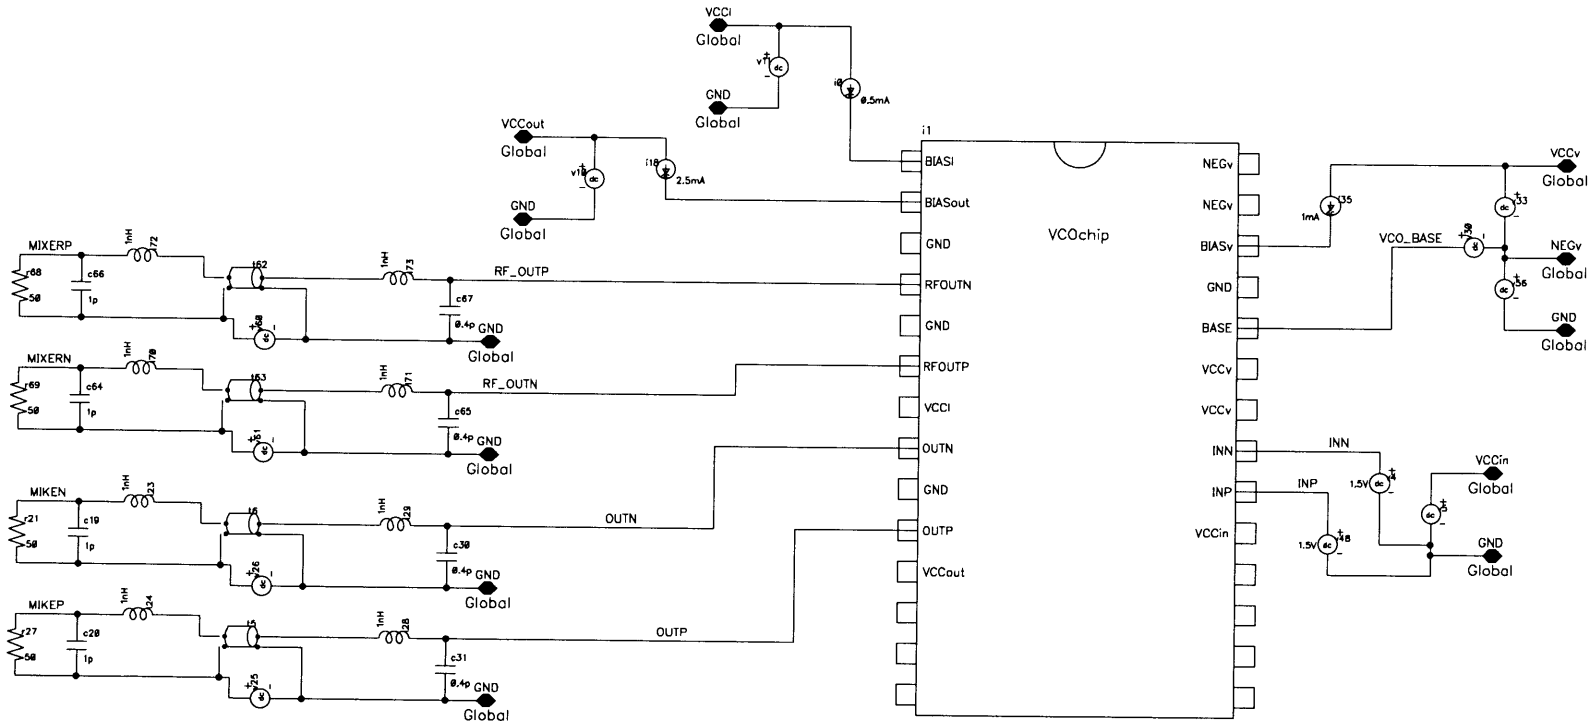
Simulations indicate that the amplitude of the steady-state oscillation is relatively unaffected until the supply is reduced to below 1 V from a nominal 1.8 V; this observation is supported by preliminary measurements. Since the noise figure of bipolar transistors is generally only weakly dependent upon the voltage across them, the consistency in oscillation amplitude should lead to a situation where the degradation occurring in the phase noise spectrum as the supply is lowered is outpaced by the power savings. However, given an available resonator quality factor, overall system noise requirements will still dictate the lower bound on power. The oscillator noise spectral density has been gauged to be -118dBc/Hz at a 1 MHz offset from a 1.7 GHz carrier for the nominal design conditions of 1.8 V, 5 mA. Oscillation is supported down to power consumption levels of 3.2 mW from a 0.91 V supply, achieving -96 dBc/Hz at the same carrier offset.

References

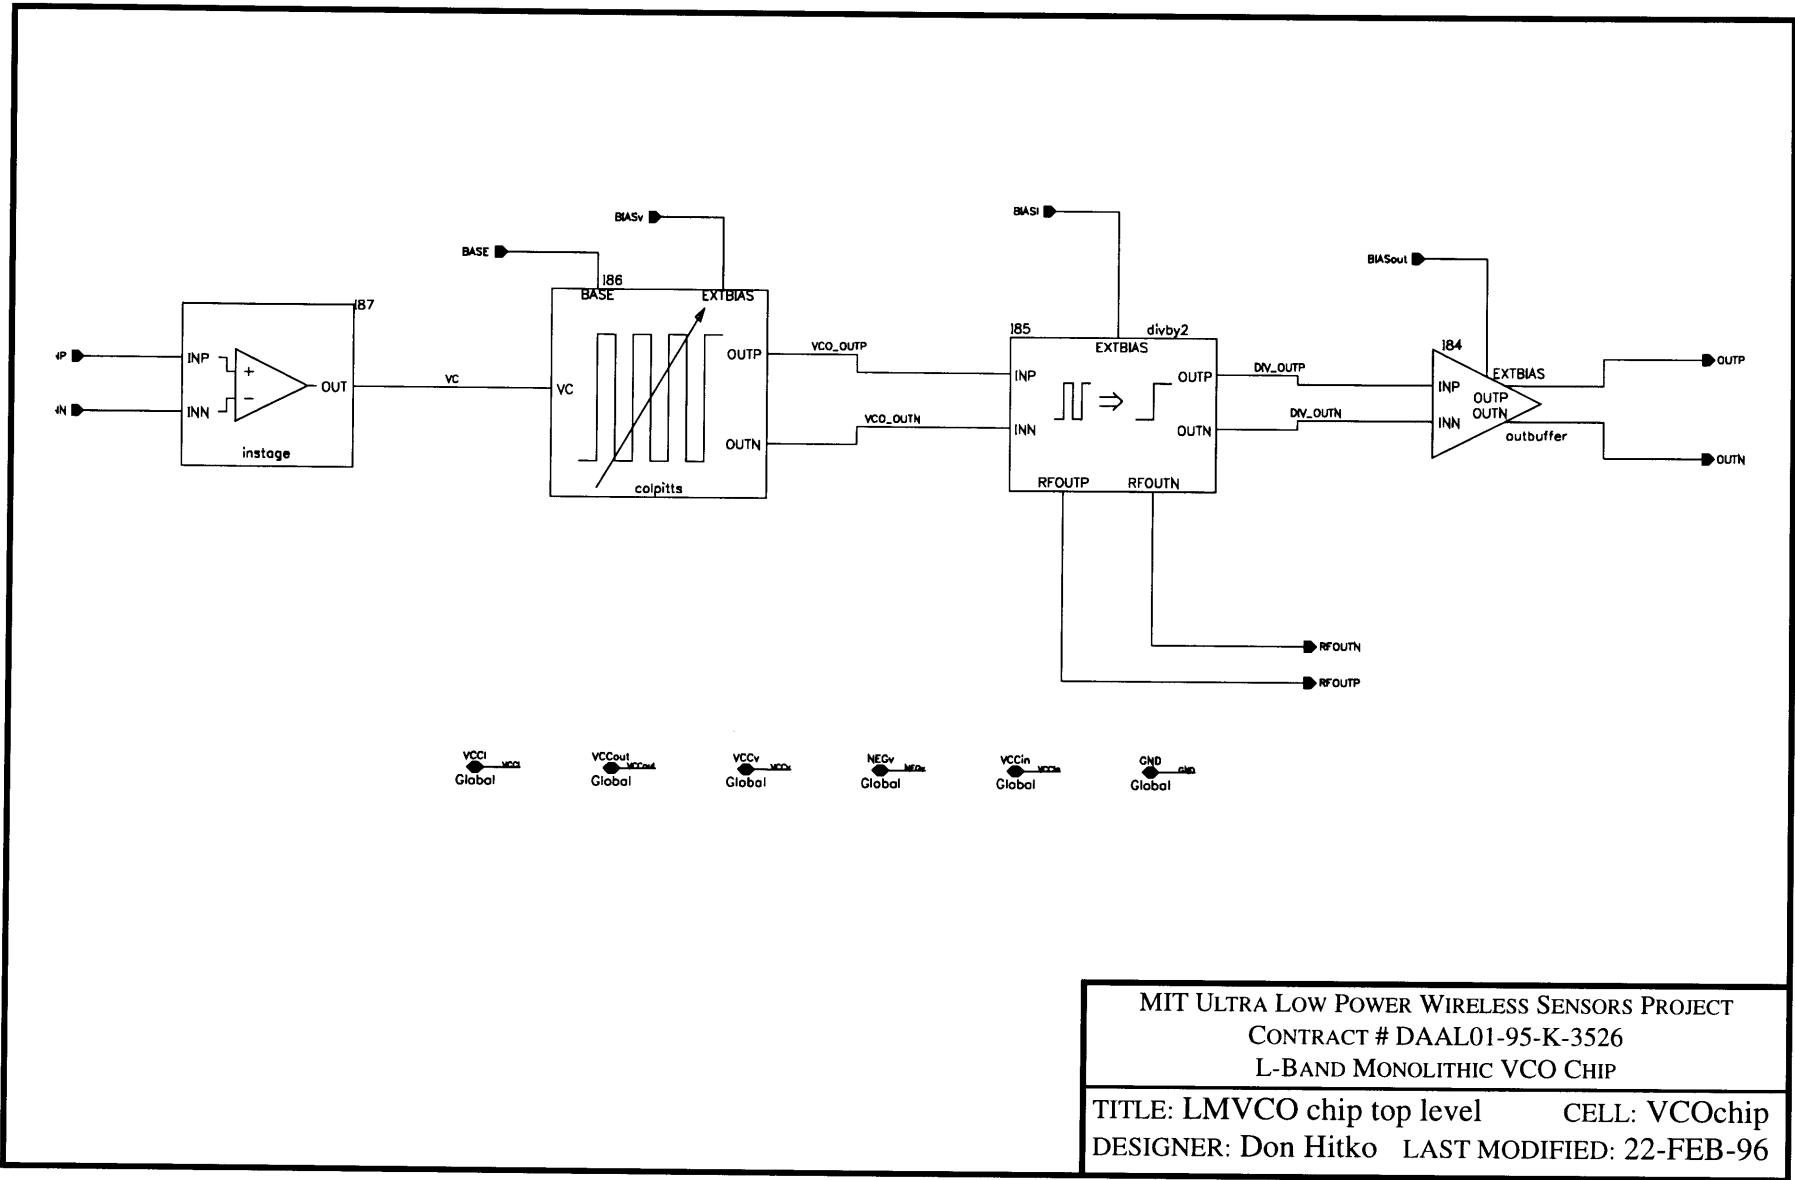
1. D. B. Leeson, "A Simple Model of Feedback Oscillator Noise Spectrum," *Proceedings of the IEEE*, February 1966, pp. 329-330.
2. H. K. Lindenmeier, "Noise Matching Techniques in Transistor Oscillators," 1980 *IEEE Circuits and Systems Symposium*, pp. 1052-1055.
3. K. Kurokawa, "Some Basic Characteristics of Broadband Negative Resistance Oscillator Circuits," *Bell System Technical Journal*, July-August 1969, pp. 1937-1955.

Appendix:

**LMVCO Chip
Schematics**

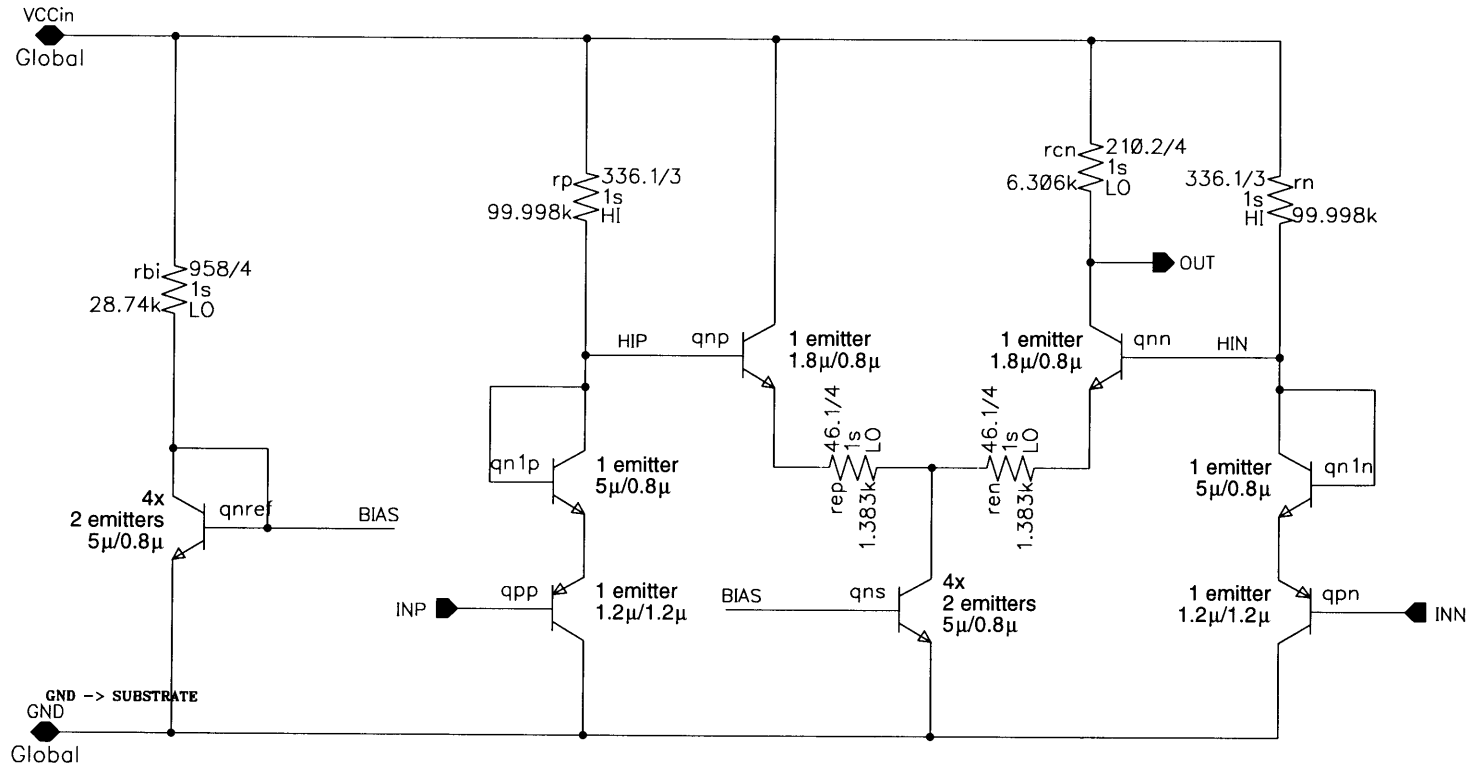


MIT ULTRA LOW POWER WIRELESS SENSORS PROJECT
CONTRACT # DAAL01-95-K-3526
L-BAND MONOLITHIC VCO CHIP
TITLE: Usage of VCO chip CELL: useVCOchip
DESIGNER: Don Hitko LAST MODIFIED: 22-FEB-96



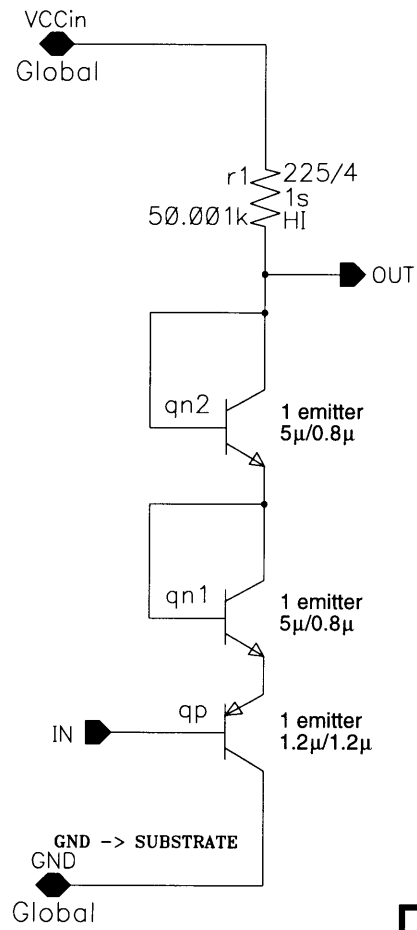
MIT ULTRA LOW POWER WIRELESS SENSORS PROJECT
 CONTRACT # DAAL01-95-K-3526
 L-BAND MONOLITHIC VCO CHIP

TITLE: LMVCO chip top level CELL: VCOchip
 DESIGNER: Don Hitko LAST MODIFIED: 22-FEB-96

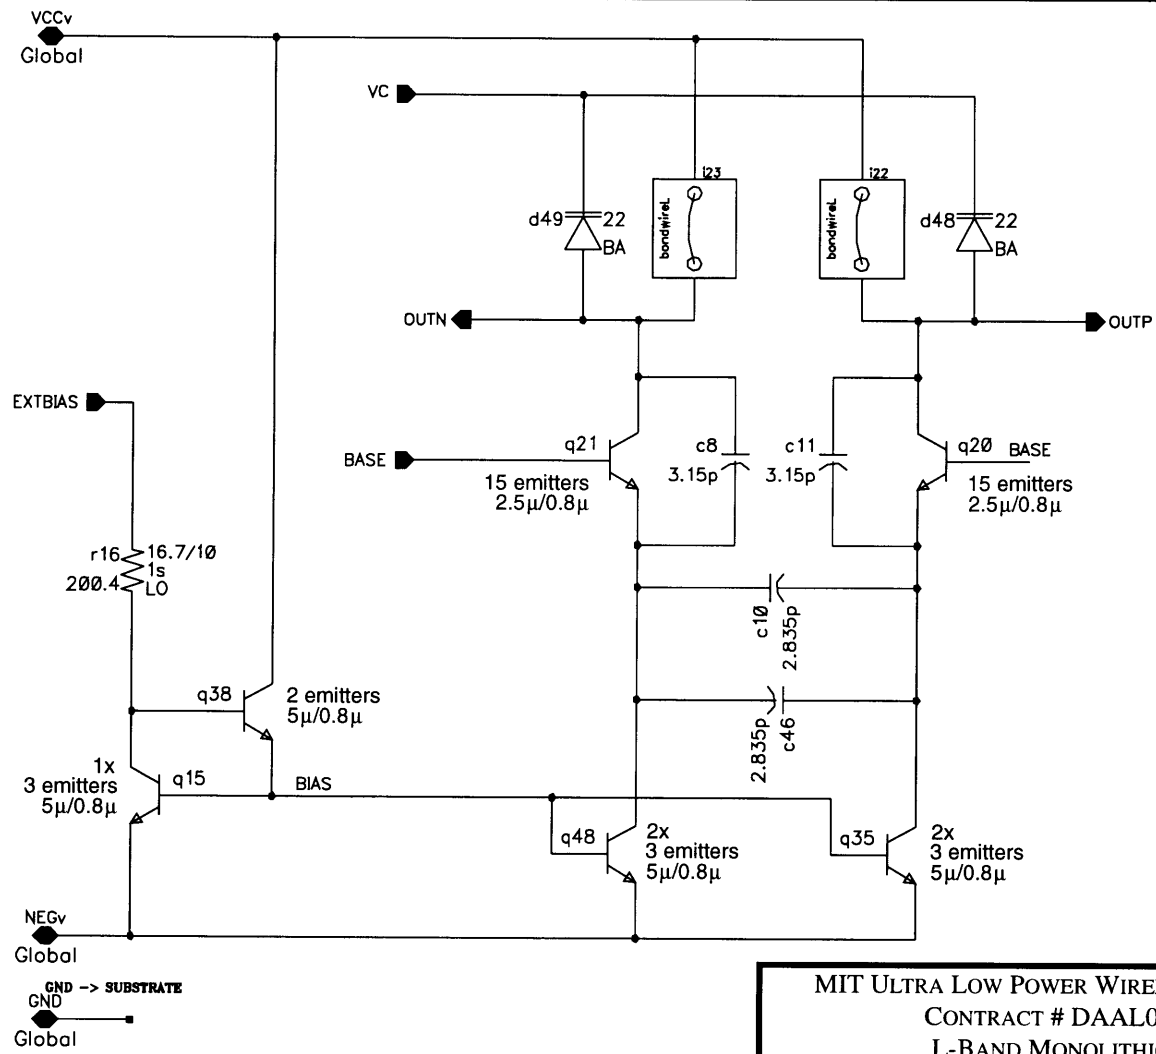


MIT ULTRA LOW POWER WIRELESS SENSORS PROJECT
 CONTRACT # DAAL01-95-K-3526
 L-BAND MONOLITHIC VCO CHIP

TITLE: Differential Input Stage CELL: instage
 DESIGNER: Don Hitko LAST MODIFIED: 22-FEB-96

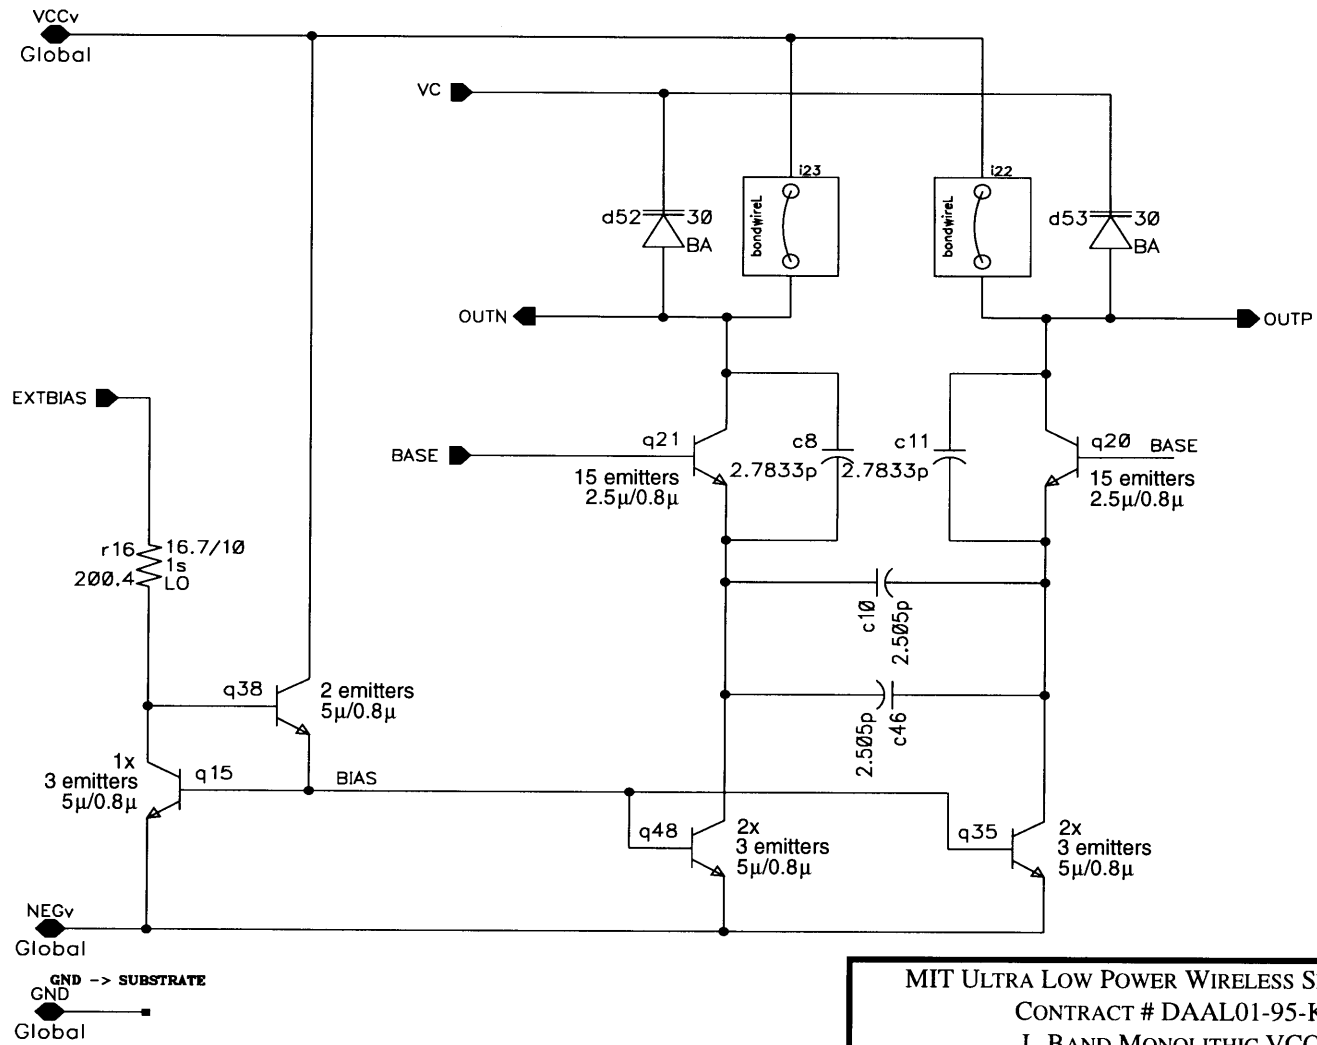


MIT ULTRA LOW POWER WIRELESS SENSORS PROJECT	
CONTRACT # DAAL01-95-K-3526	
L-BAND MONOLITHIC VCO CHIP	
TITLE: Input Level Shifting Stage	CELL: inshift
DESIGNER: Don Hitko	LAST MODIFIED: 22-FEB-96



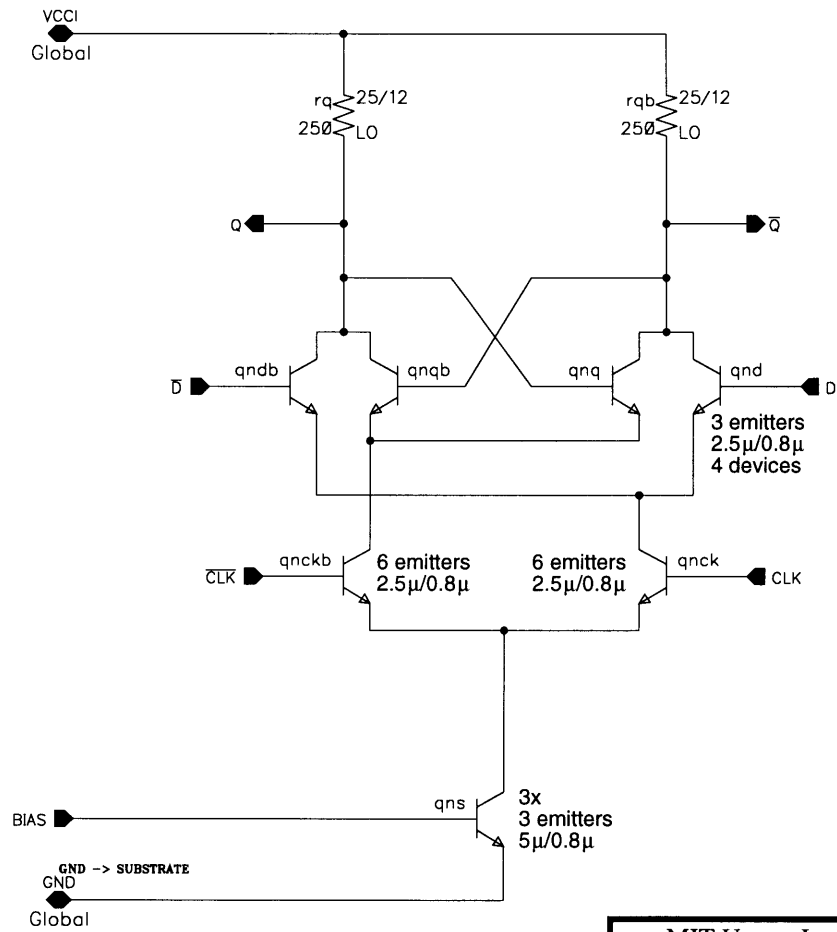
NEGv
Global
GND -> SUBSTRATE
GND
Global

MIT ULTRA LOW POWER WIRELESS SENSORS PROJECT
 CONTRACT # DAAL01-95-K-3526
 L-BAND MONOLITHIC VCO CHIP
 TITLE: Differential Colpitts VCO CELL: colpitts
 DESIGNER: Don Hitko LAST MODIFIED: 22-FEB-96

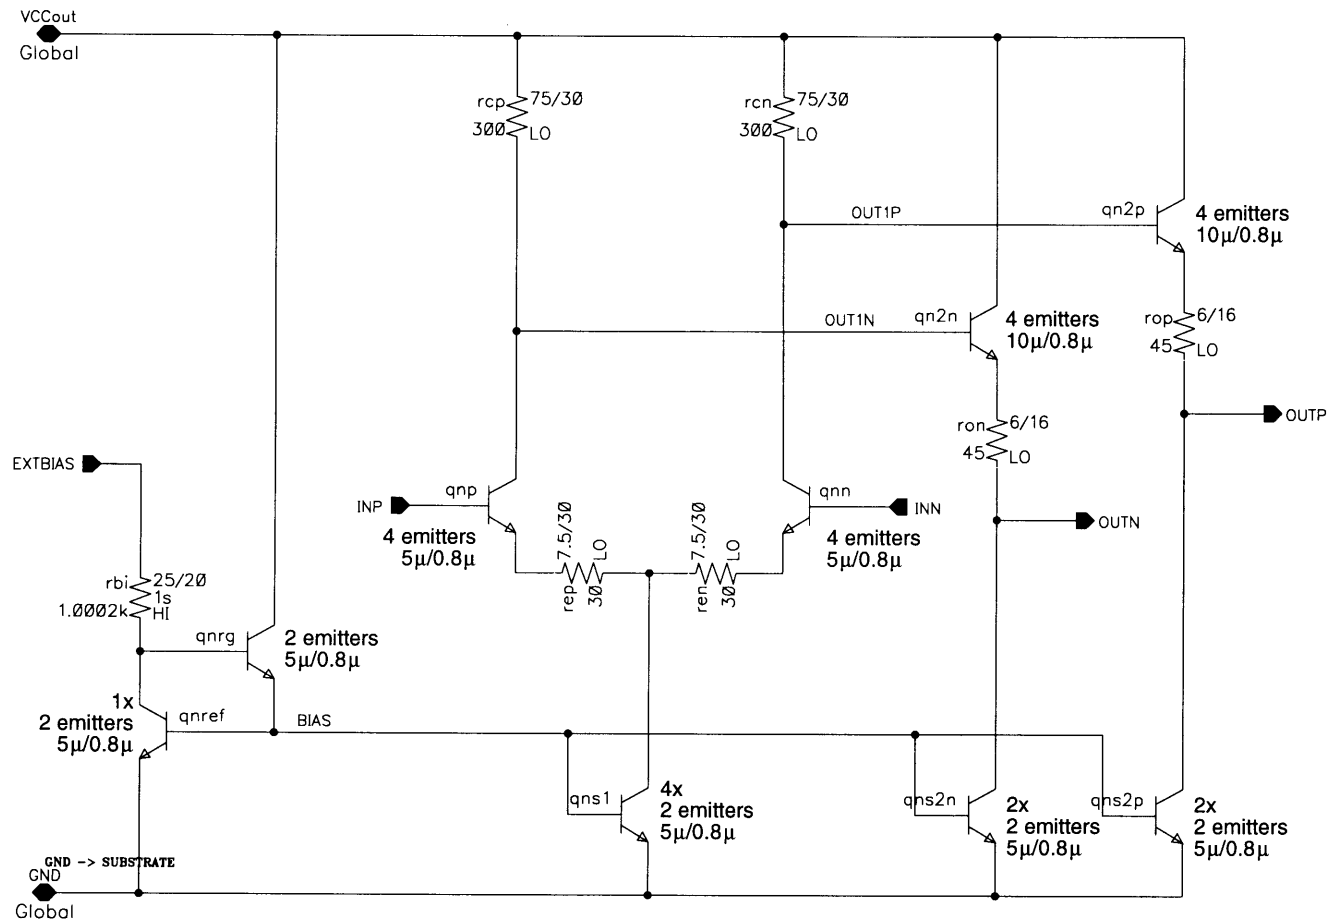


MIT ULTRA LOW POWER WIRELESS SENSORS PROJECT
 CONTRACT # DAAL01-95-K-3526
 L-BAND MONOLITHIC VCO CHIP

TITLE: 1 Input Diff. Colpitts VCO CELL: colpitts1
 DESIGNER: Don Hitko LAST MODIFIED: 22-FEB-96



MIT ULTRA LOW POWER WIRELESS SENSORS PROJECT CONTRACT # DAAL01-95-K-3526 L-BAND MONOLITHIC VCO CHIP	
TITLE: D Flip-Flop	CELL: dff
DESIGNER: Don Hitko LAST MODIFIED: 22-FEB-96	



MIT ULTRA LOW POWER WIRELESS SENSORS PROJECT
 CONTRACT # DAAL01-95-K-3526
 L-BAND MONOLITHIC VCO CHIP

TITLE: 50Ω Output Driver CELL: outbuffer
 DESIGNER: Don Hitko LAST MODIFIED: 22-FEB-96